arXiv:1611.05162v2 [cs.LG] 17 Nov 2016

Net-Trim: A Layer-wise Convex Pruning of Deep Neural Networks

Alireza Aghasi, Nam Nguyen and Justin Romberg*

Abstract

Model reduction is a highly desirable process for deep neural networks. While large networks are theoretically capable of learning arbitrarily complex models, overfitting and model redundancy negatively affects the prediction accuracy and model variance. Net-Trim is a layer-wise convex framework to prune (sparsify) deep neural networks. The method is applicable to neural networks operating with the rectified linear unit (ReLU) as the nonlinear activation. The basic idea is to retrain the network layer by layer keeping the layer inputs and outputs close to the originally trained model, while seeking a sparse transform matrix. We present both the parallel and cascade versions of the algorithm. While the former enjoys computational distributability, the latter is capable of achieving simpler models. In both cases, we mathematically show a consistency between the retrained model and the initial trained network. We also derive the general sufficient conditions for the recovery of a sparse transform matrix. In the case of standard Gaussian training samples of dimension N being fed to a layer, and s being the maximum number of nonzero terms across all columns of the transform matrix, we show that $\mathcal{O}(s \log N)$ samples are enough to accurately learn the layer model.

1 Introduction

In the context of universal approximation, neural networks can represent functions of arbitrary complexity when the network is equipped with sufficiently large number of layers and neurons [16]. Such model flexibility has made the artificial deep neural network a pioneer machine learning tool over the past decades (see [19] for a comprehensive review of deep networks). Basically, given unlimited training data and computational resources, deep neural networks are able to learn arbitrarily complex data models.

*AA and NN are with the Department of Mathematical Sciences, IBM T.J. Watson Research Center, Yorktown Heights, NY. JR is with the Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA.
Contact: aaghasi@us.ibm.com

In practice, the capability of collecting huge amount of data is often restricted. Thus, learning complicated networks with millions of parameters from limited training data can easily lead to the overfitting problem. Over the past years, various methods have been proposed to reduce overfitting via regularizing techniques and pruning strategies [18, 13, 20, 23]. However, the complex and non-convex behavior of the underlying model barricades the use of theoretical tools to analyze the performance of such techniques.

In this paper, we present an optimization framework, namely *Net-Trim*, which is a layer-wise convex scheme to sparsify deep neural networks. The proposed framework can be viewed from both theoretical and computational standpoints. Technically speaking, each layer of a neural network consists of an affine transformation (to be learned by the data) followed by a nonlinear unit. The nested composition of such mappings forms a highly nonlinear model, learning which requires optimizing a complex and non-convex objective. Net-Trim applies to a network which is already trained. The basic idea is to reduce the network complexity layer by layer, assuring that each layer response stays close to the initial trained network.

More specifically, the training data is transmitted through the learned network layer by layer. Within each layer we propose an optimization scheme which promotes weight sparsity, while enforcing a consistency between the resulting response and the trained network response. In a sense, if we consider each layer response to the transmitted data as a checkpoint, Net-Trim assures the checkpoints remain roughly the same, while a simpler path between the checkpoints is discovered. A favorable leverage of Net-Trim is the possibility of convex formulation, when the ReLU is employed as the nonlinear unit across the network.

Figure 1 demonstrates the pruning capability of Net-Trim for a sample network. The neural network used for this example classifies 200 points positioned on the 2D plane into two separate classes based on their label. The points within each class lie on nested spirals to classify which we use a neural network with two hidden layers of each 200 neurons (the reader is referred to the Experiments section for more technical details). Figures 1(a), (b) present the weighted adjacency matrix and partial network topology, relating the hidden layers before and after retraining. With only a negligible change to the overall network response, Net-Trim is able to prune more than 93% of the links among the neurons, and bring a significant model reduction to the problem. Even when the neural network is trained using sparsifying weight regularizers (here, dropout [20] and ℓ_1 penalty), application of the Net-Trim yields a major additional reduction in the model complexity as illustrated in Figures 1(c) and 1(d).

Net-Trim is particularly useful when the number of training samples is limited. While overfitting is likely to occur in such scenarios, Net-Trim reduces the complexity of the model by setting a significant portion of weights at each layer to zero, yet maintaining a similar relationship between the input data and network response. This capability can also be viewed from a different perspective, that Net-Trim simplifies the process of determining the network size. In other words, the network used at the training phase can be oversized and present more degrees of freedom than what the data require. Net-Trim would automatically reduce the network size to an order matching the data.

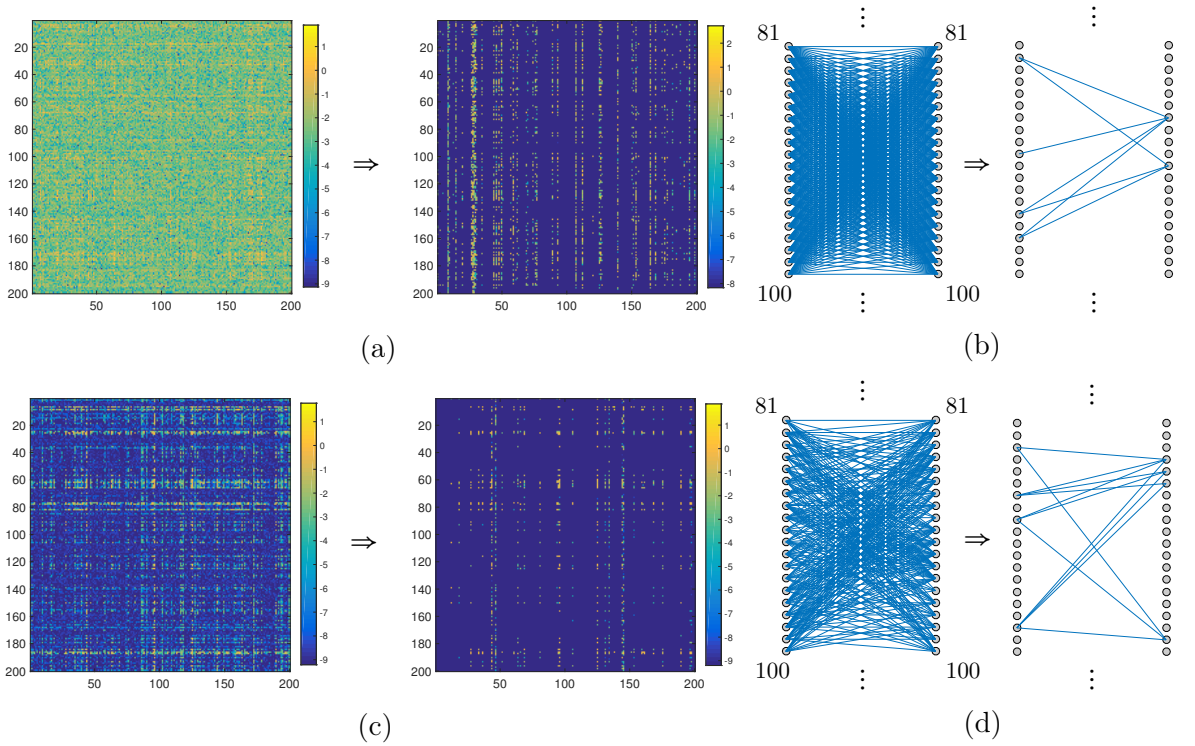

Figure 1: Net-Trim pruning performance on classification of points within nested spirals; (a) left: the weighted adjacency matrix relating the two hidden layers after training; right: the adjacency matrix after the application of Net-Trim causing more than 93% of the weights to vanish; (b) partial network topology relating neurons 81 to 100 of the hidden layers, before and after retraining; (c) left: the adjacency matrix after training the network with dropout and ℓ_1 regularization; right: Net-Trim is yet able to find a model which is over 7 times sparser than the model on the left; (d) partial network topology before and after retraining for panel (c)

Finally, a favorable property of Net-Trim is its post-processing nature. It simply processes the network layer-wise responses regardless of the training strategy used to build the model. Hence, the proposed framework can be easily combined with the state-of-the-art training techniques for deep neural networks.

1.1 Previous Work

In the recent years there has been increasing interest in the mathematical analysis of deep networks. These efforts are mainly in the context of characterizing the minimizers of the underlying cost function. In [4], the authors show that some deep forward networks can be learned accurately in polynomial time, assuming that all the edges of the network have random

weights. They propose a layer-wise algorithm where the weights are recovered sequentially at each layer. In [17], Kawaguchi establishes an exciting result showing that regardless of being highly non-convex, the square loss function of a deep neural network inherits interesting geometric structures. In particular, under some independence assumptions, all the local minima are also the global ones. In addition, the saddle points of the loss function possess special properties which guide the optimization algorithms to avoid them.

The geometry of the loss function is also studied in [10], where the authors bring a connection between spin-glass models in physics and fully connected neural networks. On the other hand, Giryes *et al.* recently provide the link between deep neural networks and compressed sensing [14], where they show that feedforward networks are able to preserve the distance of the data at each layer by using tools from compressed sensing. There are other works on formulating the training of feedforward networks as an optimization problem [7, 6, 5]. The majority of cited works approach to understand the neural networks by sequentially studying individual layers, which is also the approach taken in this paper.

On the more practical side, one of the notorious issues with training a complicated deep neural network concerns overfitting when the amount of training data is limited. There have been several approaches trying to address this issue, among those are the use of regularizations such as ℓ_1 and ℓ_2 penalty [18, 13]. These methods incorporate a penalty term to the loss function to reduce the complexity of the underlying model. Due to the non-convex nature of the underlying problem, mathematically characterizing the behavior of such regularizers is an impossible task and most literature in this area are based on heuristics. Another approach is to apply the early stopping of training as soon as the performance of a validation set starts to get worse.

More recently, a new way of regularization is proposed by Hinton et. al. [20] called Dropout. It involves temporarily dropping a portion of hidden activations during the training. In particular, for each sample, roughly 50% of the activations are randomly removed on the forward pass and the weights associated with these units are not updated on the backward pass. Combining all the examples will result in a representation of a huge ensemble of neural networks, which offers excellent generalization capability. Experimental results on several tasks indicate that Dropout frequently and significantly improves the classification performance of deep architectures. More recently, LeCun *et al.* proposed an extension of Dropout named DropConnect [23]. It is essentially similar to Dropout, except that it randomly removes the connections rather than the activations. As a result, the procedure introduces a dynamic sparsity on the weights.

The aforementioned regularization techniques (e.g ℓ_1, Dropout, and DropConnect) can be seen as methods to sparsify the neural network, which result in reducing the model complexity. Here, sparsity is understood as either reducing the connections between the nodes or decreasing the number of activations. Beside avoiding overfitting, computational favor of sparse models is preferred in applications where quick predictions are required.

1.2 Summary of the Technical Contributions

Our post-training scheme applies multiple convex programs with the ℓ_1 cost to prune the weight matrices on different layers of the network. Formally, we denote $\boldsymbol{Y}^{\ell-1}$ and $\boldsymbol{Y}^{\ell}$ as the given input and output of the ℓ-th layer of the trained neural network, respectively. The mapping between the input and output of this layer is performed via the weight matrix $\boldsymbol{W}_\ell$ and the nonlinear activation unit σ: $\boldsymbol{Y}^{\ell} = \sigma(\boldsymbol{W}_\ell^T \boldsymbol{Y}^{\ell-1})$. To perform the pruning at this layer, Net-Trim focuses on addressing the following optimization:

$$\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \left\|\sigma\left(\boldsymbol{U}^T \boldsymbol{Y}^{\ell-1}\right) - \boldsymbol{Y}^{\ell}\right\|_F \leq \epsilon, \tag{1}$$

where ϵ is a user-specific parameter that controls the consistence of the output $\boldsymbol{Y}^{\ell}$ before and after retraining and $\|\boldsymbol{U}\|_1$ is the sum of absolute entries of $\boldsymbol{U}$, which is essentially the ℓ_1 norm to enforce sparsity on $\boldsymbol{U}$.

When $\sigma(.)$ is taken to be the ReLU, we are able to provide a convex relaxation to (1). We will show that $\hat{\boldsymbol{W}}_\ell$, the solution to the underlying convex program, is not only sparser than $\boldsymbol{W}_\ell$, but the error accumulated over the layers due to the ϵ-approximation constraint will not significantly explode. In particular in Theorem 1 we show that if $\hat{\boldsymbol{Y}}^{\ell}$ is the ℓ-th layer after retraining, i.e., $\hat{\boldsymbol{Y}}^{\ell} = \sigma(\hat{\boldsymbol{W}}_\ell{}^T \hat{\boldsymbol{Y}}^{\ell-1})$, then for a network with normalized weight matrices

$$\left\|\hat{\boldsymbol{Y}}^{\ell} - \boldsymbol{Y}^{\ell}\right\|_F \leq \ell\epsilon.$$

Basically, the error propagated by the Net-Trim at any layer is at most a multiple of ϵ. This property suggests that the network constructed by the Net-Trim is sparser, while capable of achieving a similar outcome. Another attractive feature of this scheme is its computational distributability, i.e., the convex programs could be solved independently.

Also, in this paper we propose a cascade version of Net-Trim, where the output of the retraining at the previous layer is fed to the next layer as the input of the optimization. In particular, we present a convex relaxation to

$$\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \left\|\sigma\left(\boldsymbol{U}^T \hat{\boldsymbol{Y}}^{\ell-1}\right) - \boldsymbol{Y}^{\ell}\right\|_F \leq \epsilon_\ell,$$

where $\hat{\boldsymbol{Y}}^{\ell-1}$ is the retrained output of the $(\ell-1)$-th layer and ϵ_ℓ has a closed form expression to maintain feasibility of the resulting program. Again, for a network with normalized weight matrices, in Theorem 2 we show that

$$\left\|\hat{\boldsymbol{Y}}^{\ell} - \boldsymbol{Y}^{\ell}\right\|_F \leq \epsilon_1 \gamma^{(\ell-1)/2}.$$

Here $\gamma > 1$ is a constant inflation rate that can be arbitrarily close to 1 and controls the magnitude of ϵ_ℓ. Because of the more adaptive pruning, cascade Net-Trim may yield sparser solutions at the expense of not being computationally parallelizable.

Finally, for redundant networks with limited training samples, we will discuss that a simpler network (in terms of sparsity) with identical performance can be explored by setting $\epsilon = 0$ in (1). We will derive general sufficient conditions for the recovery of such sparse model via the proposed convex program. As an insightful case, we show that when a layer is probed with standard Gaussian samples (e.g., applicable to the first layer), learning the simple model can be performed with much fewer samples than the layer degrees of freedom. More specifically, consider $\boldsymbol{X} \in \mathbb{R}^{N\times P}$ to be a Gaussian matrix, where each column represents an input sample, and $\boldsymbol{W}$ a sparse matrix, with at most s nonzero terms on each column, from which the layer response is generated, i.e., $\boldsymbol{Y} = \sigma(\boldsymbol{W}^T\boldsymbol{X})$. In Theorem 3 we state that when $P = \mathcal{O}(s \log N)$, with overwhelming probability, $\boldsymbol{W}$ can be accurately learned through the proposed convex program.

As will be detailed, the underlying analysis steps beyond the standard measure concentration arguments used in the compressed sensing literature (cf. §8 in [12]). We contribute by establishing concentration inequalities for the sum of dependent random matrices.

1.3 Notations and Organization of the Paper

The remainder of the paper is structured as follows. In Section 2, we formally present the network model used in the paper. The proposed pruning schemes, both the parallel and cascade Net-Trim are presented and discussed in Section 3. The material includes insights on developing the algorithms and detailed discussions on the consistency of the retraining schemes. Section 4 is devoted to the convex analysis of the proposed framework. We derive the unique optimality conditions for the recovery of a sparse weight matrix through the proposed convex program. We then use this tool to derive the number of samples required for learning a sparse model in a Gaussian sample setup. In Section 5 we report some retraining experiments and the improvement that Net-Trim brings to the model reduction and robustness. Finally, Section 6 presents some discussions on extending the Net-Trim framework, future outlines and concluding remarks.

As a summary of the notations, our presentation mainly relies on multidimensional calculus. We use bold characters to denote vectors and matrices. Considering a matrix $\boldsymbol{A}$ and the index sets Γ_1, and Γ_2, we use $\boldsymbol{A}_{\Gamma_1,:}$ to denote the matrix obtained by restricting the rows of $\boldsymbol{A}$ to Γ_1. Similarly, $\boldsymbol{A}_{:,\Gamma_2}$ denotes the restriction of $\boldsymbol{A}$ to the columns specified by Γ_2, and $\boldsymbol{A}_{\Gamma_1,\Gamma_2}$ is the submatrix with the rows and columns restricted to Γ_1 and Γ_2, respectively.

Given a matrix $\boldsymbol{X} = [x_{m,n}] \in \mathbb{R}^{M\times N}$, we use $\|\boldsymbol{X}\|_1 \triangleq \sum_{m=1}^{M}\sum_{n=1}^{N}|x_{m,n}|$ to denote the sum of matrix absolute entries[1] and $\|\boldsymbol{X}\|_F$ to denote the Frobenius norm. For a given vector $\boldsymbol{x}$, $\|\boldsymbol{x}\|_0$ denotes the cardinality of $\boldsymbol{x}$, supp $\boldsymbol{x}$ denotes the set of indices with non-zero entries from $\boldsymbol{x}$, and supp^c $\boldsymbol{x}$ is the complement set. Mainly in the proofs, we use the notation $\boldsymbol{x}^+$ to denote $\max(\boldsymbol{x}, 0)$. The $\max(., 0)$ operation applied to a vector or matrix acts on every component

[1]The formal induced norm $\|\boldsymbol{X}\|_1$ has a different definition, however, for a simpler formulation we use a similar notation

individually. finally, following the MATLAB convention, the vertical concatenation of two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ (i.e., $[\boldsymbol{a}^T, \boldsymbol{b}^T]^T$) is sometimes denoted by $[\boldsymbol{a}; \boldsymbol{b}]$ in the text.

2 Feed Forward Network Model

In this section, we introduce some notational conventions related to a feed forward network model, which will be used frequently in the paper. Considering a feed forward neural network, we assume to have P training samples $\boldsymbol{x}_p$, $p = 1, \cdots, P$, where $\boldsymbol{x}_p \in \mathbb{R}^N$ is an input to the network. We stack up the samples in a matrix $\boldsymbol{X} \in \mathbb{R}^{N \times P}$, structured as

$$\boldsymbol{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_P].$$

The final output of the network is denoted by $\boldsymbol{Z} \in \mathbb{R}^{M \times P}$, where each column $\boldsymbol{z}_p \in \mathbb{R}^M$ of $\boldsymbol{Z}$ is a response to the corresponding training column $\boldsymbol{x}_p$ in $\boldsymbol{X}$. We consider a network with L layers, where the activations are taken to be rectified linear units. Associated with each layer ℓ, we have a weight matrix $\boldsymbol{W}_\ell$ such that

$$\boldsymbol{Y}^{(\ell)} = \max\left(\boldsymbol{W}_\ell^T \boldsymbol{Y}^{(\ell-1)}, 0\right), \qquad \ell = 1, \cdots, L, \tag{2}$$

and

$$\boldsymbol{Y}^{(0)} = \boldsymbol{X}, \quad \boldsymbol{Y}^{(L)} = \boldsymbol{Z}. \tag{3}$$

Basically, the outcome of the ℓ-th layer is $\boldsymbol{Y}^{(\ell)} \in \mathbb{R}^{N_\ell \times P}$, which is generated by applying the adjoint of $\boldsymbol{W}_\ell \in \mathbb{R}^{N_{\ell-1} \times N_\ell}$ to $\boldsymbol{Y}^{(\ell-1)}$ and going through a component-wise $\max(., 0)$ operation. Clearly in this setup $N_0 = N$ and $N_L = M$. A trained neural network as outlined in (2) and (3) is represented by $\mathcal{TN}(\{\boldsymbol{W}_\ell\}_{\ell=1}^L, \boldsymbol{X})$. Figure 2(a) demonstrates the architecture of the proposed network.

For the sake of theoretical analysis, throughout the paper we focus on networks with normalized weights as follows.

Definition 1. *A given neural network $\mathcal{TN}(\{\boldsymbol{W}_\ell\}_{\ell=1}^L, \boldsymbol{X})$ is link-normalized when $\|\boldsymbol{W}_\ell\|_1 = 1$ for every layer $\ell = 1, \cdots, L$.*

A general network in the form of (2) can be converted to its link-normalized version by replacing $\boldsymbol{W}_\ell$ with $\boldsymbol{W}_\ell / \|\boldsymbol{W}_\ell\|_1$, and $\boldsymbol{Y}^{(\ell+1)}$ with $\boldsymbol{Y}^{(\ell+1)} / \prod_{j=0}^{\ell} \|\boldsymbol{W}_j\|_1$. Since $\max(\alpha x, 0) = \alpha \max(x, 0)$ for $\alpha > 0$, any weight processing on a network of the form (2) can be applied to the link-normalized version and later transferred to the original domain via a suitable scaling. Subsequently, all the results presented in this paper are stated for a link-normalized network.

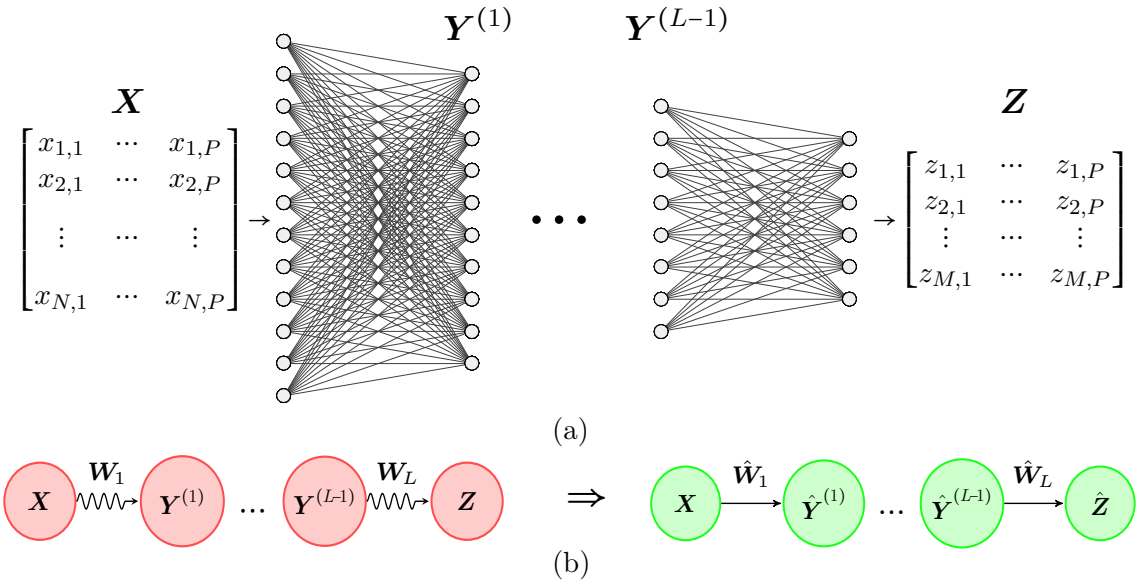

Figure 2: (a) Network architecture and notations; (b) the main retraining idea: keeping the layer outcomes close to the initial trained model while finding a simpler path relating each layer input to the output

3 Pruning the Network

Our pruning strategy relies on redesigning the network so that for the same training data each layer outcomes stay more or less close to the initial trained model, while the weights associated with each layer are replaced with sparser versions to reduce the model complexity. Figure 2(b) presents the main idea, where the complex paths between the layer outcomes are replaced with simple paths.

Consider the first layer, where $\boldsymbol{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_P]$ is the layer input, $\boldsymbol{W} = [\boldsymbol{w}_1, \cdots, \boldsymbol{w}_M]$ the layer coefficient matrix, and $\boldsymbol{Y} = [y_{m,p}]$ the layer outcome. We require the new coefficient matrix $\hat{\boldsymbol{W}}$ to be sparse and the new response to be close to $\boldsymbol{Y}$. Using the sum of absolute entries as a proxy to promote sparsity, a natural strategy to retrain the layer is addressing the nonlinear program

$$\hat{\boldsymbol{W}} = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \left\|\max\left(\boldsymbol{U}^T\boldsymbol{X}, 0\right) - \boldsymbol{Y}\right\|_F \leq \epsilon. \tag{4}$$

Despite the convex objective, the constraint set in (4) is non-convex. However, we may approximate it with a convex set by imposing $\boldsymbol{Y}$ and $\hat{\boldsymbol{Y}} = \max(\hat{\boldsymbol{W}}^T\boldsymbol{X}, 0)$ to have similar activation patterns. More specifically, knowing that $y_{m,p}$ is either zero or positive, we enforce the $\max(., 0)$ argument to be negative when $y_{m,p} = 0$, and close to $y_{m,p}$ elsewhere. To present

the convex formulation, for $\boldsymbol{V} = [v_{m,p}]$ we use the notation

$$\begin{cases} \sum\limits_{m,p:\, y_{m,p}>0} \left(\boldsymbol{u}_m^T \boldsymbol{x}_p - y_{m,p}\right)^2 \le \epsilon^2 \\ \boldsymbol{u}_m^T \boldsymbol{x}_p \le v_{m,p} \qquad\qquad for\ \ m,p:\ y_{m,p}=0 \end{cases} \iff \quad \boldsymbol{U} \in \mathcal{C}_\epsilon(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{V}). \tag{5}$$

Based on this definition, a convex proxy to (4) is

$$\hat{\boldsymbol{W}} = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_\epsilon(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{0}). \tag{6}$$

Basically, depending on the value of $y_{m,p}$, a different constraint is imposed on $\boldsymbol{u}_m^T \boldsymbol{x}_p$ to emulate the ReLU operation. For a simpler formulation throughout the paper, we use a similar notation as $\mathcal{C}_\epsilon(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{V})$ for any constraints of the form (5) parametrized by given $\boldsymbol{X}$, $\boldsymbol{Y}$, $\boldsymbol{V}$ and ϵ.

As a first step towards establishing a retraining framework applicable to the entire network, we show that the solution of (6) satisfies the constraint in (4) and the outcome of the retrained layer stays controllably close to $\boldsymbol{Y}$.

Proposition 1. *Let $\hat{\boldsymbol{W}}$ be the solution to (6). For $\hat{\boldsymbol{Y}} = \max(\hat{\boldsymbol{W}}^T \boldsymbol{X}, 0)$ being the retrained layer response, $\|\hat{\boldsymbol{Y}} - \boldsymbol{Y}\|_F \le \epsilon$.*

Based on the above exploratory, we propose two schemes to retrain the neural network; one explores a computationally distributable nature and the other proposes a cascading scheme to retrain the layers sequentially. The general idea which originates from the relaxation in (6) is referred to as the Net-Trim, specified by the parallel or cascade nature.

3.1 Parallel Net-Trim

The parallel Net-Trim is a straightforward application of the convex program (6) to each layer in the network. Basically, each layer is processed independently based on the initial model input and output, without taking into account the retraining result from the previous layer. Specifically, denoting $\boldsymbol{Y}^{(\ell-1)}$ and $\boldsymbol{Y}^{(\ell)}$ as the input and output of the ℓ-th layer of the initially trained neural network (see equation (2)), we propose to relearn the coefficient matrix $\boldsymbol{W}_\ell$ via the convex program

$$\hat{\boldsymbol{W}}_\ell = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_\epsilon\left(\boldsymbol{Y}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{0}\right). \tag{7}$$

The optimization (7) can be independently applied to every layer in the network and hence computationally distributable. The pseudocode for the parallel Net-Trim is presented as Algorithm 1.

With reference to the constraint in (7), if we only retrain the ℓ-th layer, the output of the retrained layer is in the ϵ-neighborhood of that before retraining. However, when all the layers

Algorithm 1 Parallel Net-Trim

Input: $\boldsymbol{X}$, $\epsilon > 0$, and link-normalized $\boldsymbol{W}_1, \cdots, \boldsymbol{W}_L$
$\boldsymbol{Y}^{(0)} \leftarrow \boldsymbol{X}$
for $\ell = 1, \cdots, L$ **do**
  $\boldsymbol{Y}^{(\ell)} \leftarrow \max\left(\boldsymbol{W}_\ell^T \boldsymbol{Y}^{(\ell-1)}, 0\right)$ `% generating link-normalized layer outcomes`
end for
for all $\ell = 1, \cdots, L$ **do**
  $\hat{\boldsymbol{W}}_\ell \leftarrow \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_\epsilon\left(\boldsymbol{Y}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{0}\right)$ `% retraining`
end for
Output: $\hat{\boldsymbol{W}}_1, \cdots, \hat{\boldsymbol{W}}_L$

are retrained through (7), an immediate question would be whether the retrained network produces an output which is controllably close to the initially trained model. In the following theorem, we show that the retrained error does not blow up across the layers and remains a multiple of ϵ.

Theorem 1. *Given a link-normalized network $\mathcal{TN}(\{\boldsymbol{W}_\ell\}_{\ell=1}^L, \boldsymbol{X})$ with layer outcomes $\boldsymbol{Y}^{(\ell)}$ as sketched in (2) and (3), consider retraining each layer individually via*

$$\hat{\boldsymbol{W}}_\ell = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_{\epsilon_\ell}\left(\boldsymbol{Y}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{0}\right). \tag{8}$$

For the retrained network $\mathcal{TN}(\{\hat{\boldsymbol{W}}_\ell\}_{\ell=1}^L, \boldsymbol{X})$ with layer outcomes $\hat{\boldsymbol{Y}}^{(\ell)} = \max(\hat{\boldsymbol{W}}_\ell^T \hat{\boldsymbol{Y}}^{(\ell-1)}, 0)$,

$$\left\|\hat{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\right\|_F \le \sum_{j=1}^{\ell} \epsilon_j. \tag{9}$$

When all the layers are retrained with a fixed parameter ϵ (as in Algorithm 1), the following corollary simply bounds the overall discrepancy.

Corollary 1. *Using Algorithm 1, the ultimate network outcome obeys*

$$\left\|\hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)}\right\|_F \le L\epsilon.$$

We would like to note that the network conversion to a link-normalized version is only for the sake of presenting the theoretical results in a more compact form. In practice such conversion is not necessary and to retrain layer ℓ we can take $\epsilon = \epsilon_r \|\boldsymbol{Y}^{(\ell)}\|_F$, where ϵ_r plays a similar role as ϵ for a link-normalized network.

3.2 Cascade Net-Trim

Unlike the parallel scheme, where each layer is retrained independently, in the cascade approach the outcome of a retrained layer is used to retrain the next layer. To better explain the mechanics, consider starting the cascade process by retraining the first layer as before, through

$$\hat{\boldsymbol{W}}_1 = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_{\epsilon_1}\left(\boldsymbol{X}, \boldsymbol{Y}^{(1)}, \boldsymbol{0}\right). \tag{10}$$

Setting $\hat{\boldsymbol{Y}}^{(1)} = \max(\hat{\boldsymbol{W}}_1^T \boldsymbol{X}, 0)$ to be the outcome of the retrained layer, to retrain the second layer, we ideally would like to address a similar program as (10) with $\hat{\boldsymbol{Y}}^{(1)}$ as the input and $\boldsymbol{Y}^{(2)}$ being the output reference, i.e.,

$$\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_{\epsilon_2}\left(\hat{\boldsymbol{Y}}^{(1)}, \boldsymbol{Y}^{(2)}, \boldsymbol{0}\right). \tag{11}$$

However, there is no guarantee that program (11) is feasible, that is, there exists a matrix $\boldsymbol{W} = [\boldsymbol{w}_1, \cdots, \boldsymbol{w}_{N_2}]$ such that

$$\begin{cases} \sum\limits_{m,p:\, y^{(2)}_{m,p}>0} \left(\boldsymbol{w}_m^T \hat{\boldsymbol{y}}_p^{(1)} - y^{(2)}_{m,p}\right)^2 \le \epsilon_2^2 & \\ \boldsymbol{w}_m^T \hat{\boldsymbol{y}}_p^{(1)} \le 0 \qquad\qquad for \ \ m,p:\ y^{(2)}_{m,p} = 0 & \end{cases}. \tag{12}$$

If instead of $\hat{\boldsymbol{Y}}^{(1)}$ the constraint set (11) was parameterized by $\boldsymbol{Y}^{(1)}$, a natural feasible point would have been $\boldsymbol{W}_2$. Now that $\hat{\boldsymbol{Y}}^{(1)}$ is a perturbed version of $\boldsymbol{Y}^{(1)}$, the constraint set needs to be slacked to maintain the feasibility of $\boldsymbol{W}_2$. In this context, one may easily verify that

$$\boldsymbol{W}_2 \in \mathcal{C}_{\epsilon_2}\left(\hat{\boldsymbol{Y}}^{(1)}, \boldsymbol{Y}^{(2)}, \boldsymbol{W}_2^T \hat{\boldsymbol{Y}}^{(1)}\right) \tag{13}$$

as long as

$$\epsilon_2^2 \ge \sum_{m,p:\, y^{(2)}_{m,p}>0} \left(\boldsymbol{w}_{2,m}^T \hat{\boldsymbol{y}}_p^{(1)} - y^{(2)}_{m,p}\right)^2, \tag{14}$$

where $\boldsymbol{w}_{2,m}$ is the m-th column of $\boldsymbol{W}_2$. Basically the constraint set in (13) is a slacked version of the constraint set in (12), where the right hand side quantities in the corresponding inequalities are sufficiently extended to maintain the feasibility of $\boldsymbol{W}_2$.

Following this line of argument, in the cascade Net-Trim we propose to retrain the first layer through (10). For every subsequent layer, $\ell = 2, \cdots, L$, the retrained weighting matrix is obtained via

$$\hat{\boldsymbol{W}}_\ell = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_{\epsilon_\ell}\left(\hat{\boldsymbol{Y}}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{W}_\ell^T \hat{\boldsymbol{Y}}^{(\ell-1)}\right), \tag{15}$$

where for $\boldsymbol{W}_\ell = [\boldsymbol{w}_{\ell,1}, \cdots, \boldsymbol{w}_{\ell,N_\ell}]$ and $\gamma_\ell \geq 1$,

$$\epsilon_\ell^2 = \gamma_\ell \sum_{m,p:\, y_{m,p}^{(\ell)} > 0} \left(\boldsymbol{w}_{\ell,m}^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - y_{m,p}^{(\ell)} \right)^2.$$

The constants $\gamma_\ell \geq 1$ (referred to as the *inflation rates*) are free parameters, which control the sparsity of the resulting matrices. After retraining the ℓ-th layer we set

$$\hat{\boldsymbol{Y}}^{(\ell)} = \max\left(\hat{\boldsymbol{W}}_\ell^T \hat{\boldsymbol{Y}}^{(\ell-1)}, 0 \right),$$

and use this outcome to retrain the next layer. Algorithm 2 presents the pseudo-code to implement the cascade Net-Trim for $\epsilon_1 = \epsilon$ and a constant inflation rate, γ, across all the layers.

Algorithm 2 Cascade Net-Trim

Input: $\boldsymbol{X}$, $\epsilon > 0$, $\gamma > 1$ and link-normalized $\boldsymbol{W}_1, \cdots, \boldsymbol{W}_L$
$\boldsymbol{Y} \leftarrow \max\left(\boldsymbol{W}_1^T \boldsymbol{X}, 0 \right)$
$\hat{\boldsymbol{W}}_1 \leftarrow \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_\epsilon(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{0})$
$\hat{\boldsymbol{Y}} \leftarrow \max(\hat{\boldsymbol{W}}_1^T \boldsymbol{X}, 0)$
for $\ell = 2, \cdots, L$ **do**
  $\boldsymbol{Y} \leftarrow \max(\boldsymbol{W}_\ell^T \boldsymbol{Y}, 0)$
  $\epsilon \leftarrow (\gamma \sum_{m,p:y_{m,p}>0} (\boldsymbol{w}_{\ell,m}^T \hat{\boldsymbol{y}}_p - y_{m,p})^2)^{1/2}$ `% ` $\boldsymbol{w}_{\ell,m}$ `is the` m`-th column of` $\boldsymbol{W}_\ell$
  $\hat{\boldsymbol{W}}_\ell \leftarrow \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_\epsilon(\hat{\boldsymbol{Y}}, \boldsymbol{Y}, \boldsymbol{W}_\ell^T \hat{\boldsymbol{Y}})$
  $\hat{\boldsymbol{Y}} \leftarrow \max(\hat{\boldsymbol{W}}_\ell^T \hat{\boldsymbol{Y}}, 0)$
end for
Output: $\hat{\boldsymbol{W}}_1, \cdots, \hat{\boldsymbol{W}}_L$

In the following theorem, we prove that the outcome of the retrained network produced by Algorithm 2 is close to that of the network before retraining.

Theorem 2. *Given a link-normalized network* $\mathcal{TN}(\{\boldsymbol{W}_\ell\}_{\ell=1}^L, \boldsymbol{X})$ *with layer outcomes* $\boldsymbol{Y}^{(\ell)}$*, consider retraining the first layer as* (10) *and the subsequent layers via* (15)*, such that*

$$\epsilon_\ell^2 = \gamma_\ell \sum_{m,p:\, y_{m,p}^{(\ell)} > 0} \left(\boldsymbol{w}_{\ell,m}^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - y_{m,p}^{(\ell)} \right)^2,$$

$\hat{\boldsymbol{Y}}^{(\ell)} = \max(\hat{\boldsymbol{W}}_\ell^T \hat{\boldsymbol{Y}}^{(\ell-1)}, 0)$*,* $\hat{\boldsymbol{Y}}^{(1)} = \max(\hat{\boldsymbol{W}}_1^T \boldsymbol{X}, 0)$ *and* $\gamma_\ell > 1$*. For* $\mathcal{TN}(\{\hat{\boldsymbol{W}}_\ell\}_{\ell=1}^L, \boldsymbol{X})$ *being the retrained network*

$$\left\| \hat{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)} \right\|_F \leq \epsilon_1 \sqrt{\prod_{j=2}^{\ell} \gamma_j}. \tag{16}$$

When $\epsilon_1 = \epsilon$ and all the layers are retrained with a fixed inflation rate (as in Algorithm 2), the following corollary of Theorem 2 bounds the network overall discrepancy.

Corollary 2. *Using Algorithm 2, the ultimate network outcome obeys*

$$\left\| \hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)} \right\|_F \leq \gamma^{\frac{(L-1)}{2}} \epsilon.$$

Similar to the parallel Net-Trim, the cascade Net-Trim can also be performed without a link-normalization by simply setting $\epsilon = \epsilon_r \|\boldsymbol{Y}^{(1)}\|_F$.

3.3 Retraining the Last Layer

Commonly, the last layer in a neural network is not subject to an activation function and a standard linear model applies, i.e., $\boldsymbol{Y}^{(L)} = \boldsymbol{W}_L^T \boldsymbol{Y}^{(L-1)}$. This linear outcome may be directly exploited for regression purposes or pass through a soft-max function to produce the scores for a classification task.

In this case, to retrain the layer we simply need to seek a sparse weight matrix under the constraint that the linear outcomes stay close before and after retraining. More specifically,

$$\hat{\boldsymbol{W}}_L = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \left\| \boldsymbol{U}^T \boldsymbol{Y}^{(L-1)} - \boldsymbol{Y}^{(L)} \right\|_F \leq \epsilon_L. \tag{17}$$

In the case of cascade Net-Trim,

$$\hat{\boldsymbol{W}}_L = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \left\| \boldsymbol{U}^T \hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L)} \right\|_F \leq \epsilon_L, \tag{18}$$

and the feasibility of the program is established for

$$\epsilon_L^2 = \gamma_L \left\| \boldsymbol{W}_L^T \hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L)} \right\|_F^2, \qquad \gamma \geq 1. \tag{19}$$

It can be shown that the results stated earlier in Theorems 1 and 2 regarding the overall discrepancy of the network generalize to a network with linear activation at the last layer.

Proposition 2. *Consider a link-normalized network $\mathcal{TN}(\{\boldsymbol{W}_\ell\}_{\ell=1}^L, \boldsymbol{X})$, where a standard linear model applies to the last layer.*

(a) If the first $L-1$ layers are retrained according to the process stated in Theorem 1 and the last layer is retrained through (17)*, then*

$$\left\| \hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)} \right\|_F \leq \sum_{\ell=1}^{L} \epsilon_j.$$

(b) If the first $L-1$ layers are retrained according to the process stated in Theorem 2 and the last layer is retrained through (18) *and* (19)*, then*

$$\left\| \hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)} \right\|_F \leq \epsilon_1 \sqrt{\prod_{j=2}^{L} \gamma_j}.$$

While the cascade Net-Trim is designed in way that infeasibility is never an issue, one can take a slight risk of infeasibility in retraining the last layer to further reduce the overall discrepancy. More specifically, if the value of ϵ_L in (18) is replaced with $\kappa\epsilon_L$ for some $\kappa \in (0,1)$, we may reduce the overall discrepancy by the factor κ, without altering the sparsity pattern of the first $L-1$ layers. It is however clear that in this case there is no guarantee that program (18) remains feasible and multiple trials may be needed to tune κ. We will refer to κ as the risk coefficient and will present some examples in Section 5, which use it as a way to control the final discrepancy in a cascade framework.

4 Convex Analysis and Model Learning

In this section we will focus on redundant networks, where the mapping between a layer input and the corresponding output can be established via various weight matrices. As an example, this could be the case when insufficient training samples are used to train a large network. We will show that in this case, if the relation between the layer input and output can be established via a sparse weight matrix, under some conditions such matrix could be uniquely identified through the core Net-Trim program in (6).

As noted above, in the case of a redundant layer, for a given input $\boldsymbol{X}$ and output $\boldsymbol{Y}$, the relation $\boldsymbol{Y} = \max(\boldsymbol{W}^T\boldsymbol{X}, 0)$ can be established via more than one $\boldsymbol{W}$. In this case we hope to find a sparse $\boldsymbol{W}$ by setting $\epsilon = 0$ in (6). For this value of ϵ our central convex program reduces to

$$\hat{\boldsymbol{W}} = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \begin{cases} \boldsymbol{u}_m^T\boldsymbol{x}_p = y_{m,p} & for \ \ m,p: \ y_{m,p} > 0 \\ \boldsymbol{u}_m^T\boldsymbol{x}_p \leq 0 & for \ \ m,p: \ y_{m,p} = 0 \end{cases},$$

which decouples into M convex programs, each searching for the m-th column in $\hat{\boldsymbol{W}}$:

$$\hat{\boldsymbol{w}}_m = \arg\min_{\boldsymbol{w}} \ \|\boldsymbol{w}\|_1 \quad s.t. \quad \begin{cases} \boldsymbol{w}^T\boldsymbol{x}_p = y_{m,p} & for \ \ p: \ y_{m,p} > 0 \\ \boldsymbol{w}^T\boldsymbol{x}_p \leq 0 & for \ \ p: \ y_{m,p} = 0 \end{cases}.$$

For a more concise representation, we drop the m index and given a vector $\boldsymbol{y} \in \mathbb{R}^P$ focus on the convex program

$$\min_{\boldsymbol{w}} \ \|\boldsymbol{w}\|_1 \quad s.t. \quad \begin{cases} \boldsymbol{X}_{:,\Omega}^T\boldsymbol{w} = \boldsymbol{y}_\Omega \\ \boldsymbol{X}_{:,\Omega^c}^T\boldsymbol{w} \leq \boldsymbol{0} \end{cases}, \quad \text{where} \quad \Omega = \{p : y_p > 0\}. \tag{20}$$

In the remainder of this section we analyze the optimality conditions for (20), and show how they can be linked to the identification of a sparse solution. The program can be cast as

$$\min_{\boldsymbol{w},\boldsymbol{s}} \ \|\boldsymbol{w}\|_1 \qquad s.t. \qquad \tilde{\boldsymbol{X}} \begin{bmatrix} \boldsymbol{w} \\ \boldsymbol{s} \end{bmatrix} = \boldsymbol{y}, \quad \boldsymbol{s} \preceq \boldsymbol{0}, \tag{21}$$

where

$$\tilde{\boldsymbol{X}} = \begin{bmatrix} \boldsymbol{X}^T_{:,\Omega} & \boldsymbol{0} \\ \boldsymbol{X}^T_{:,\Omega^c} & -\boldsymbol{I} \end{bmatrix} \quad \text{and} \quad \boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_\Omega \\ \boldsymbol{0} \end{bmatrix}.$$

For a general $\tilde{\boldsymbol{X}}$, not necessarily structured as above, the following result states the sufficient conditions under which a sparse pair $(\boldsymbol{w}^*, \boldsymbol{s}^*)$ is the unique minimizer to (21).

Proposition 3. *Consider a pair $(\boldsymbol{w}^*, \boldsymbol{s}^*) \in (\mathbb{R}^{n_1}, \mathbb{R}^{n_2})$, which is feasible for the convex program (21). If there exists a vector $\boldsymbol{\Lambda} = [\Lambda_\ell] \in \mathbb{R}^{n_1+n_2}$ in the range of $\tilde{\boldsymbol{X}}^T$ with entries satisfying*

$$\begin{cases} -1 < \Lambda_\ell < 1 & \ell \in supp^c\, \boldsymbol{w}^* \\ 0 < \Lambda_{n_1+\ell} & \ell \in supp^c\, \boldsymbol{s}^* \end{cases}, \qquad \begin{cases} \Lambda_\ell = sign(w^*_\ell) & \ell \in supp\, \boldsymbol{w}^* \\ \Lambda_{n_1+\ell} = 0 & \ell \in supp\, \boldsymbol{s}^* \end{cases}, \tag{22}$$

and for $\tilde{\Gamma} = supp\, \boldsymbol{w}^* \cup \{n_1 + supp\, \boldsymbol{s}^*\}$ the restricted matrix $\tilde{\boldsymbol{X}}_{:,\tilde{\Gamma}}$ is full column rank, then the pair $(\boldsymbol{w}^*, \boldsymbol{s}^*)$ is the unique solution to (21).

The proposed optimality result can be related to the unique identification of a sparse $\boldsymbol{w}^*$ from rectified observations of the form $\boldsymbol{y} = \max(\boldsymbol{X}^T\boldsymbol{w}^*, 0)$. Clearly, the structure of the feature matrix $\boldsymbol{X}$ plays the key role here, and the construction of the dual certificate stated in Proposition 3 entirely relies on that. As an insightful case, we show that when $\boldsymbol{X}$ is a *Gaussian matrix* (that is, the elements of $\boldsymbol{X}$ are i.i.d values drawn from a standard normal distribution), learning $\boldsymbol{w}^*$ can be performed with much fewer samples than the layer degrees of freedom.

Theorem 3. *Let $\boldsymbol{w}^* \in \mathbb{R}^N$ be an arbitrary s-sparse vector, $\boldsymbol{X} \in \mathbb{R}^{N\times P}$ a Gaussian matrix representing the samples and $\mu > 1$ a fixed value. Given $P = (6s+1)\mu \log N$ observations of the type $\boldsymbol{y} = \max(\boldsymbol{X}^T\boldsymbol{w}^*, 0)$, with probability exceeding $1 - N^{1-\mu}$ the vector $\boldsymbol{w}^*$ can be learned exactly through* (20).

The standard Gaussian assumption for the feature matrix $\boldsymbol{X}$ allows us to relate the number of training samples to the number of active links in a layer. Such feature structure could be a realistic assumption for the first layer of the neural network. As shown in the proof of Theorem 3, because of the dependence of the set Ω to the entries in $\boldsymbol{X}$, the standard concentration of measure framework for independent random matrices is not applicable here. Instead, we will need to establish concentration bounds for the sum of dependent random matrices.

Because of the contribution the weight matrices have to the distribution of $\boldsymbol{Y}^{(1)}, \cdots, \boldsymbol{Y}^{(L)}$, without restrictive assumptions, a similar type of analysis for the subsequent layers seems significantly harder and left as a possible future work. Yet, Theorem 3 is a good reference for the number of required training samples to learn a sparse model for Gaussian (or approximately Gaussian) samples. While we focused on each decoupled problem individually, for observations of the type $\boldsymbol{Y} = \max(\boldsymbol{W}^{*T}\boldsymbol{X}, 0)$, using the union bound, an exact identification of $\boldsymbol{W}^*$ can be warranted as a corollary of Theorem 3.

Corollary 3. *Consider an arbitrary matrix $\boldsymbol{W}^* = [\boldsymbol{w}_1^*, \cdots, \boldsymbol{w}_M^*] \in \mathbb{R}^{N\times M}$, where $s_m = \|\boldsymbol{w}_m^*\|_0$, and $0 < s_{\min} \leq s_m \leq s_{\max}$ for $m = 1, \cdots, M$. For $\boldsymbol{X} \in \mathbb{R}^{N\times P}$ being a Gaussian matrix, set $\boldsymbol{Y} = \max(\boldsymbol{W}^{*T}\boldsymbol{X}, 0)$. Then if*

$$\mu > (1 + \log_N M)\, \frac{6s_{\max} + 1}{6s_{\min} + 1}$$

and $P = (6s_{\max}+1)\mu \log N$, for $\epsilon = 0$, $\boldsymbol{W}^*$ can be accurately learned through (6) with probability exceeding

$$1 - \sum_{m=1}^{M} N^{1-\mu\frac{6s_m+1}{6s_{\max}+1}}.$$

4.1 Pruning Partially Clustered Neurons

As discussed above, in the case of $\epsilon = 0$, program (6) decouples into M smaller convex problems, which could be addressed individually and computationally cheaper. Clearly, for $\epsilon \neq 0$ a similar decoupling does not produce the formal minimizer to (6), but such suboptimal solution may yet significantly contribute to the pruning of the layer.

Basically, in retraining $\boldsymbol{W} \in \mathbb{R}^{N\times M}$ corresponding to a large layer with a large number of training samples, one may consider partitioning the output nodes into N_c clusters $C_1, C_2, \cdots, C_{N_c}$ such that $\cup_{k=1}^{N_c} C_k = \{1, 2, \cdots, M\}$, and solve an individual version of (6) for each cluster focusing on the underlying target nodes:

$$\tilde{\boldsymbol{W}}_{:,C_k} = \arg\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \boldsymbol{U} \in \mathcal{C}_{\epsilon_k}(\boldsymbol{X}, \boldsymbol{Y}_{C_k,:}, \boldsymbol{0}). \tag{23}$$

Solving (23) for each cluster provides the retrained submatrix associated with that cluster. The values of ϵ_k in (23) are selected in a way that ultimately the overall layer discrepancy is upper-bounded by ϵ. In this regard, a natural choice would be

$$\epsilon_k = \epsilon\sqrt{\frac{|C_k|}{M}}.$$

While $\hat{\boldsymbol{W}}$, acquired through (6), and $\tilde{\boldsymbol{W}}$ are only identical in the case of $\epsilon = 0$, the idea of clustering the output neurons into multiple groups and retraining each sublayer individually

can significantly help with breaking down large problems into computationally tractable ones. Some examples of Net-Trim with partially clustered neurons (PCN) will be presented in the experiments section. Clearly, the most distributable case is choosing a single neuron for each partition (i.e., $C_k = \{k\}$ and $N_c = M$), which results in the smallest sublayers to retrain.

4.2 Retraining Reformulation as a Quadratically Constrained Program

As discussed earlier, Net-Trim implementation requires addressing optimizations of the form

$$\min_{\boldsymbol{U}} \ \|\boldsymbol{U}\|_1 \quad s.t. \quad \begin{cases} \sum\limits_{m,p:\, y_{m,p}>0} \left(\boldsymbol{u}_m^T \boldsymbol{x}_p - y_{m,p}\right)^2 \le \epsilon^2 \\ \boldsymbol{u}_m^T \boldsymbol{x}_p \le v_{m,p} \qquad\qquad for\ \ m,p:\ y_{m,p}=0 \end{cases}, \tag{24}$$

where $\boldsymbol{U} = [\boldsymbol{u}_1, \cdots, \boldsymbol{u}_M] \in \mathbb{R}^{N\times M}$, $\boldsymbol{X} = [\boldsymbol{x}_1, \cdots, \boldsymbol{x}_P] \in \mathbb{R}^{N\times P}$, $\boldsymbol{Y} = [y_{m,p}] \in \mathbb{R}^{M\times P}$ and $\boldsymbol{V} = [v_{m,p}] \in \mathbb{R}^{M\times P}$. By the construction of the problem, all elements of $\boldsymbol{Y}$ are non-negative. In this section we represent (24) in a matrix form, which can be fed into standard quadratically constrained solvers. For this purpose we try to rewrite (24) in terms of

$$\boldsymbol{u} = vec(\boldsymbol{U}) \in \mathbb{R}^{MN},$$

where the $vec(.)$ operator converts $\boldsymbol{U}$ into a vector of length MN by stacking its columns on top of one another. Also corresponding to the subscript index sets $\{(m,p) : y_{m,p} > 0\}$ and $\{(m,p) : y_{m,p} = 0\}$ we define the complement *linear* index sets

$$\Omega = \{(m-1)P + p : y_{m,p} > 0\}, \quad \Omega^c = \{(m-1)P + p : y_{m,p} = 0\}.$$

Denoting $\boldsymbol{I}_M$ as the identity matrix of size $M \times M$, using basic properties of the Kronecker product it is straightforward to verify that

$$\boldsymbol{u}_m^T \boldsymbol{x}_p = \boldsymbol{u}^T \left(\boldsymbol{I}_M \otimes \boldsymbol{X}\right)_{:,(m-1)P+p}.$$

Basically, $\boldsymbol{u}_m^T \boldsymbol{x}_p$ is the inner product between $vec(\boldsymbol{U})$ and column $(m-1)P + p$ of $\boldsymbol{I}_M \otimes \boldsymbol{X}$. Subsequently, denoting $\boldsymbol{y} = vec(\boldsymbol{Y}^T)$ and $\boldsymbol{v} = vec(\boldsymbol{V}^T)$, we can rewrite (24) in terms of $\boldsymbol{u}$ as

$$\min_{\boldsymbol{u}} \ \|\boldsymbol{u}\|_1 \quad s.t. \quad \begin{cases} \boldsymbol{u}^T \boldsymbol{Q} \boldsymbol{u} + 2\boldsymbol{q}^T \boldsymbol{u} \le \tilde{\epsilon} \\ \boldsymbol{P}\boldsymbol{u} \le \boldsymbol{c} \end{cases}, \tag{25}$$

where

$$\boldsymbol{Q} = \left(\boldsymbol{I}_M \otimes \boldsymbol{X}\right)_{:,\Omega} \left(\boldsymbol{I}_M \otimes \boldsymbol{X}\right)_{:,\Omega}^T, \quad \boldsymbol{q} = -\left(\boldsymbol{I}_M \otimes \boldsymbol{X}\right)_{:,\Omega} \boldsymbol{y}_\Omega, \quad \tilde{\epsilon} = \epsilon^2 - \boldsymbol{y}_\Omega^T \boldsymbol{y}_\Omega \tag{26}$$

and

$$\boldsymbol{P} = \left(\boldsymbol{I}_M \otimes \boldsymbol{X}\right)_{:,\Omega^c}^T, \quad \boldsymbol{c} = \boldsymbol{v}_{\Omega^c}. \tag{27}$$

Using the formulation above allows us to cast (24) as (25), where the unknown is a vector instead of a matrix.

We can apply an additional change of variable to make (25) adaptable to standard quadratically constrained convex solvers. For this purpose we define a new vector $\tilde{\boldsymbol{u}} = [\boldsymbol{u}^+; -\boldsymbol{u}^-]$, where $\boldsymbol{u}^- = \min(\boldsymbol{u}, 0)$. This variable change naturally yields

$$\boldsymbol{u} = [\boldsymbol{I}, -\boldsymbol{I}]\tilde{\boldsymbol{u}}, \quad \|\boldsymbol{u}\|_1 = \mathbf{1}^T\tilde{\boldsymbol{u}}.$$

The convex program (25) is now cast as the quadratic program

$$\min_{\tilde{\boldsymbol{u}}} \ \mathbf{1}^T\tilde{\boldsymbol{u}} \quad s.t. \quad \begin{cases} \tilde{\boldsymbol{u}}^T\tilde{\boldsymbol{Q}}\tilde{\boldsymbol{u}} + 2\tilde{\boldsymbol{q}}^T\tilde{\boldsymbol{u}} \le \tilde{\epsilon} \\ \tilde{\boldsymbol{P}}\tilde{\boldsymbol{u}} \le \boldsymbol{c} \\ \tilde{\boldsymbol{u}} \ge \mathbf{0} \end{cases}, \tag{28}$$

where

$$\tilde{\boldsymbol{Q}} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \otimes \boldsymbol{Q}, \qquad \tilde{\boldsymbol{q}} = \begin{bmatrix} \boldsymbol{q} \\ -\boldsymbol{q} \end{bmatrix}, \qquad \tilde{\boldsymbol{P}} = \begin{bmatrix} \boldsymbol{P} & -\boldsymbol{P} \end{bmatrix}.$$

Once $\tilde{\boldsymbol{u}}^*$, the solution to (28) is found, we can obtain $\boldsymbol{u}^*$ (the solution to (25)) through the relation $\boldsymbol{u}^* = [\boldsymbol{I}, -\boldsymbol{I}]\tilde{\boldsymbol{u}}^*$. Reshaping $\boldsymbol{u}^*$ to a matrix of size $N \times M$ ultimately returns $\boldsymbol{U}^*$, the matrix solution to (24).

5 Experiments

In this section we present some learning and retraining examples to highlight the performance of the Net-Trim. To train our networks we use the *H2O* package for deep learning [3, 8], which is equipped with the well-known pruning and regularizing tools such as the dropout and ℓ_1-penalty. While developing a specialized solver to address the Net-Trim core convex program remains a future work, standard quadratic solvers such as the *IBM ILOG CPLEX* [11] and *Gurobi* [15] can be employed to address the underlying quadratic programs.

5.1 Basic Classification: Data Points on Nested Spirals

For a better demonstration of the Net-trim performance in terms of model reduction, mean accuracy and cascade vs. parallel retraining frameworks, here we focus on a low dimensional dataset. We specifically look into the classification of two set of points lying on nested spirals as shown in Figure 3(a). The dataset is embedded into the H2O package and publicly available along with the module.

As an initial experiment, we consider a network of size $2{\cdot}200{\cdot}200{\cdot}2$, which indicates the use of two hidden layers of 200 neurons each, i.e., $\boldsymbol{W}_1 \in \mathbb{R}^{2\times 200}$, $\boldsymbol{W}_2 \in \mathbb{R}^{200\times 200}$ and $\boldsymbol{W}_3 \in \mathbb{R}^{200\times 2}$. After training the model, a contour plot of the soft-max outcome, indicating the classifier, is

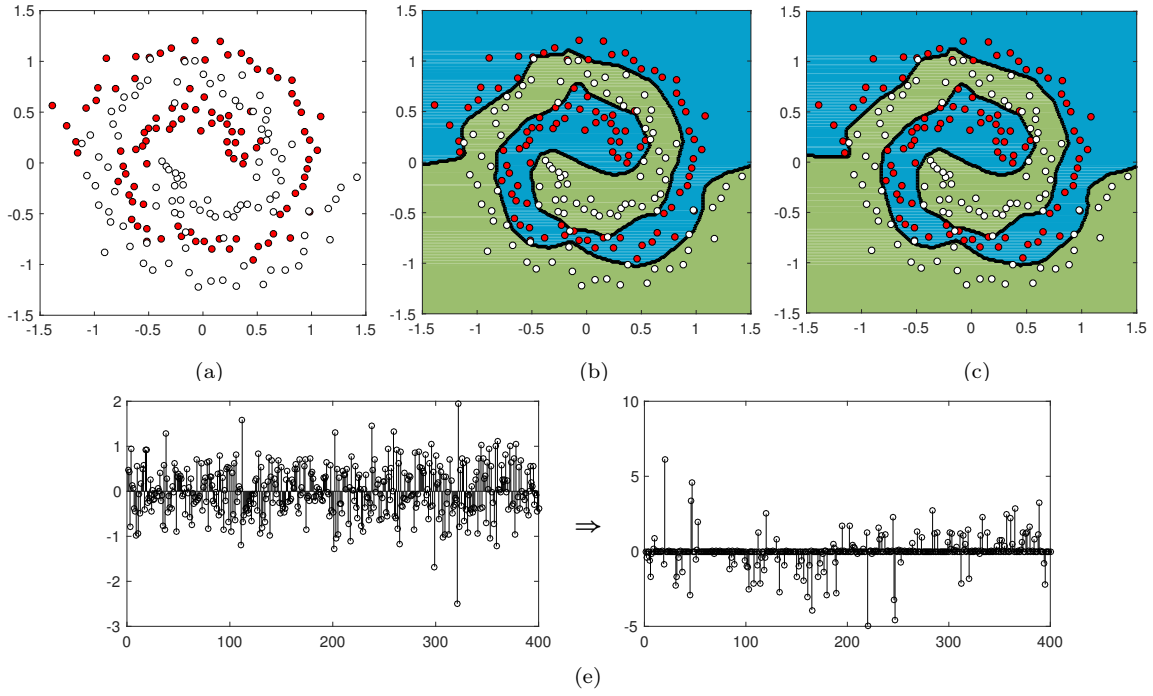

Figure 3: Classifying two set of data points on nested spirals; (a) the points corresponding to each class with different colors; (b) the soft-max contour (0.5 level-set) representing the neural net classifier; (c) the classifier after applying the Net-Trim (d) a plot of the network weights corresponding to the last layer, before (on the left side) and after (on the right side) retraining

depicted in Figure 3(b). We apply the cascade Net-Trim for $\epsilon = 0.01 \times \|\boldsymbol{Y}^{(1)}\|_F$ (the network is not link normalized), $\gamma = 1.1$ and the final risk coefficient $\kappa = 0.35$. To evaluate the difference between the network output before and after retraining, we define the relative discrepancy

$$\epsilon_{rd} = \frac{\|\boldsymbol{Z} - \hat{\boldsymbol{Z}}\|_F}{\|\boldsymbol{Z}\|_F}, \tag{29}$$

where $\boldsymbol{Z} = \boldsymbol{W}_3\boldsymbol{Y}^{(2)}$ and $\hat{\boldsymbol{Z}} = \hat{\boldsymbol{W}}_3\hat{\boldsymbol{Y}}^{(2)}$ are the network outcomes before the soft-max operation. In this case $\epsilon_{rd} = 0.046$. The classifier after retraining is presented in Figure 3(c), which shows minor difference with the original classifier in panel (b). The number of nonzero elements in $\boldsymbol{W}_1, \boldsymbol{W}_2$ and $\boldsymbol{W}_3$ are 397, 39770 and 399, respectively. After retraining, the active entries in $\hat{\boldsymbol{W}}_1, \hat{\boldsymbol{W}}_2$ and $\hat{\boldsymbol{W}}_3$ reduce to 362, 2663 and 131 elements, respectively. Basically, at the expense of a slight model discrepancy, a significant reduction in the model complexity is achieved. Figures 1(a) and 3(e) compare the cardinalities of the second and third layer weights before and after retraining.

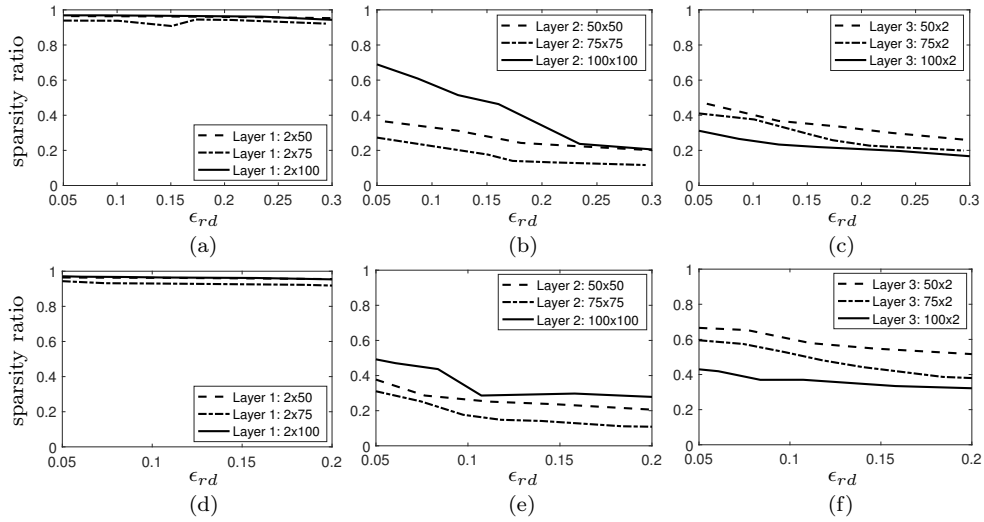

Figure 4: Sparsity ratio as a function of overall network relative mismatch for the cascade (first row) and parallel (second row) schemes

As a second experiment, we train the neural network with dropout and ℓ_1 penalty to produce a readily simple model. The number of nonzero elements in $\boldsymbol{W}_1, \boldsymbol{W}_2$ and $\boldsymbol{W}_3$ turn out to be 319, 6554 and 304, respectively. Using a similar ϵ as the first experiment, we apply the cascade Net-Trim, which produces a retrained model with $\epsilon_{rd} = 0.0183$ (the classifiers are visually identical and not shown here). The number of active entries in $\hat{\boldsymbol{W}}_1, \hat{\boldsymbol{W}}_2$ and $\hat{\boldsymbol{W}}_3$ are 189, 929 and 84, respectively. Despite the use of model reduction tools (dropout and ℓ_1 penalty) in the training phase, the Net-Trim yet zeros out a large portion of the weights in the retraining phase. The second layer weight-matrix densities before and after retraining are visually comparable in Figure 1(c).

We next perform a more extensive experiment to evaluate the performance of the cascade Net-Trim against the parallel version. Using the spiral data, we train three networks each with two hidden layers of sizes $50 \cdot 50$, $75 \cdot 75$ and $100 \cdot 100$. For the parallel retraining, we fix a value of ϵ, retrain each model 20 times and record the mean layer sparsity across these experiments (the averaging among 20 experiments is to remove the bias of local minima in the training phase). A similar process is repeated for the cascade case, where we consistently use $\gamma = 1.1$ and $\kappa = 1$. We can sweep the values of ϵ in a range to generate a class of curves relating the network relative discrepancy to each layer mean sparsity ratio, as presented in Figure 4. Here, sparsity ratio refers to the ratio of active elements to the total number of elements in the weight matrix.

A natural observation from the decreasing curves is allowing more discrepancy leads to

more level of sparsity. We also observe that for a constant discrepancy ϵ_{rd}, the cascade Net-Trim is capable of generating rather sparser networks. The contrast in sparsity is more apparent in the third layer (panel (c) vs. panel (f)). We would like to note that using $\kappa < 1$ makes the contrast even more tangible, however for the purpose of long-run simulations, here we chose $\kappa = 1$ to avoid any possible infeasibility interruptions. Finally, an interesting observation is the rather dense retrained matrices associated with the first layer. Apparently, less pruning takes place at the first layer to maximally bring the information and data structure into the network.

In Table 1 we have listed some retraining scenarios for networks of different sizes trained with dropout. Across all the experiments, we have used the cascade Net-Trim to retrain the networks and chosen ϵ small enough to warrant an overall relative discrepancy below 0.02. On the right side of the table, the number of active elements for each layer is reported, which indicates the significant model reduction for a negligible discrepancy.

Table 1: Number of active elements within each layer, before and after Net-Trim for a network trained with Dropout

	Trained Network			Net-Trim Retrained Network		
Network Size	Layer 1	Layer 2	Layer 3	Layer 1	Layer 2	Layer 3
$2 \cdot 50 \cdot 50 \cdot 2$	99	2483	100	98	467	54
$2 \cdot 75 \cdot 75 \cdot 2$	149	5594	150	149	710	72
$2 \cdot 125 \cdot 125 \cdot 2$	250	15529	250	247	3477	96
$2 \cdot 175 \cdot 175 \cdot 2$	349	30395	350	348	1743	116
$2 \cdot 200 \cdot 200 \cdot 2$	400	39668	399	399	1991	113

Table 2: Number of active elements within each layer, before and after Net-Trim for a network trained with Dropout and an ℓ_1-penalty

	Trained Network			Net-Trim Retrained Network		
Network Size	Layer 1	Layer 2	Layer 3	Layer 1	Layer 2	Layer 3
$2 \cdot 50 \cdot 50 \cdot 2$	58	1604	95	54	342	46
$2 \cdot 75 \cdot 75 \cdot 2$	96	2867	135	90	651	62
$2 \cdot 125 \cdot 125 \cdot 2$	126	5316	226	95	751	60
$2 \cdot 175 \cdot 175 \cdot 2$	171	9580	320	136	906	61
$2 \cdot 200 \cdot 200 \cdot 2$	134	8700	382	109	606	70

Table 2 reports another set of sample experiments, where dropout and ℓ_1 penalty are simultaneously employed in the training phase to prune the network. Going through a similar cascade retraining, while keeping ϵ_{rd} below 0.02, we have reported the level of additional

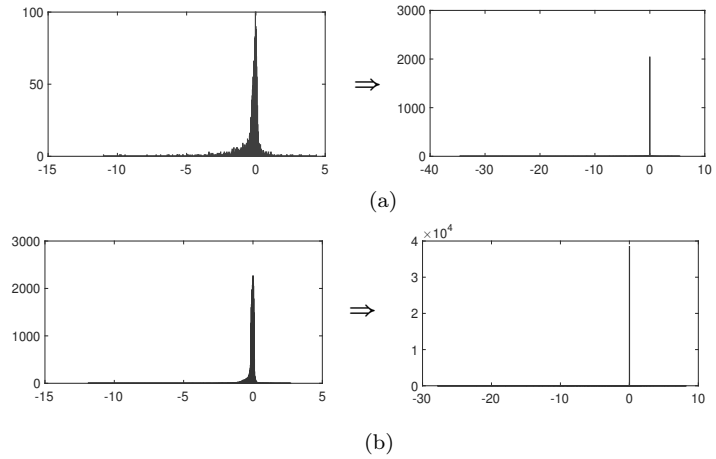

Figure 5: The weight histogram of the middle layer before and after retraining; (a) the middle layer histogram of a $2 \cdot 50 \cdot 50 \cdot 2$ network trained with dropout (left) vs. the histogram after Net-Trim (right); (b) similar plots as panel (a) for a $2 \cdot 200 \cdot 200 \cdot 2$ network

model reduction that can be achieved. Basically, the Net-Trim post processing module uses the trained model (regardless of how it is trained) to further reduce its complexity. A comparison of the network weight histograms before and after retraining may better highlight the Net-Trim performance. Figure 5 compares the middle layer weight histograms for a pair of experiments reported in Table 1.

5.2 Character Recognition

In this section we apply Net-Trim to the problem of classifying hand-written digits. For this purpose we use a fraction of the mixed national institute of standards and technology (MNIST) dataset. The set contains 60,000 training samples and 10,000 test instances. This classification problem has been well-studied in the literature, and error rates of almost 0.2% have been achieved using the full training set [23]. However, here we focus on the problem of training the models with limited samples (a fraction of the data) and show how the Net-Trim stabilizes this process.

For this problem we apply the parallel Net-Trim. Also, in order to maximally distribute the problem among independent computational nodes, we use the PCN Net-Trim with one target node. Basically, at every layer, an individual Net-Trim program is solved for every output neuron. We consider four experiments with 200, 300, 500 and 1000 training samples. For every sample set of size P, we simply select the first P rows of the original MNIST training set. Similarly, our test corresponds to the first 1000 samples of the original test set.

For the experiments with 200 and 300 training samples we use two hidden layers of each 1024 neurons. For the cases of 500 and 1000 samples, three hidden layers of each 1024 neurons

Table 3: A summary of the network architecture, sparsity ratio (SR) and classification accuracy (CA) for various training sample sizes

Sample Size (P)	200		300		500		1000	
Hidden Units	$1024 \cdot 1024$		$1024 \cdot 1024$		$1024 \cdot 1024 \cdot 1024$		$1024 \cdot 1024 \cdot 1024$	
Retraining	Before	After	Before	After	Before	After	Before	After
SR (Layer 1)	0.68	0.19	0.71	0.26	0.73	0.39	0.76	0.57
SR (Layer 2)	0.98	0.16	0.98	0.17	0.98	0.24	0.98	0.27
SR (Layer 3)	0.99	0.18	0.99	0.26	0.98	0.20	0.98	0.29
SR (Layer 4)	–	–	–	–	0.99	0.31	0.99	0.43
CA	77.5	77.7	82.2	82.6	86.1	86.5	89.2	89.8
CA (5% noise)	62.7	70.0	75.7	78.9	78.8	80.5	62.1	74.2
CA (10% noise)	46.1	55.0	61.5	70.4	62.1	65.5	39.2	52.9

are used. Table 3 summarizes the test accuracies after the initial training.

To retrain the perceptron corresponding to neuron m of layer ℓ, we use $\epsilon = \epsilon_r \|\boldsymbol{Y}_{m,:}^{(\ell)}\|$, where $\epsilon_r \in \{0.002, 0.005, 0.01\}$. With such individual neuron retraining the layer response would also obey $\|\boldsymbol{Y}^{(\ell)} - \hat{\boldsymbol{Y}}^{(\ell)}\|_F \leq \epsilon_r \|\boldsymbol{Y}^{(\ell)}\|_F$. We used three ϵ_r values to run a simple cross-validation test and find a combination of the retrained layers which produces a better network in terms of accuracy and complexity. For the 200 and 300-sample networks $\epsilon_r = 0.005$ across all three layers produced the best networks. For the 1000-sample network using $\epsilon_r = 0.01$ across all the layers, and for the 500-sample network the ϵ_r sequence $(0.005, 0.01, 0.005, 0.01)$ for the layers one through four seemed the better combinations.

Table 3 also reports the layer-wise sparsity ratio before and after retraining. We can observe the substantial improvement in sparsity ratio gained after the retraining. To see how this reduction in the model helps with the classification accuracy, we report the identification success rate for the test data as well as the cases of data being contaminated with 5% and 10% Gaussian noise. We can see that the pruning scheme improves the accuracy especially when the data are noisy. Basically, as expected by reducing the model complexity the network becomes more robust to the outliers and noisy samples.

For a deeper study of the improvement Net-Trim brings to handling noisy data, in Figure 6 we have plotted the classification accuracy against the noise level for the four trained networks. The Net-Trim improvement in accuracy becomes more noticeable as the noise level in the data increases. Finally, Figure 7 reports a selected set of the test samples, where the original network prediction was false and reducing the network complexity through the Net-Trim corrected the prediction.

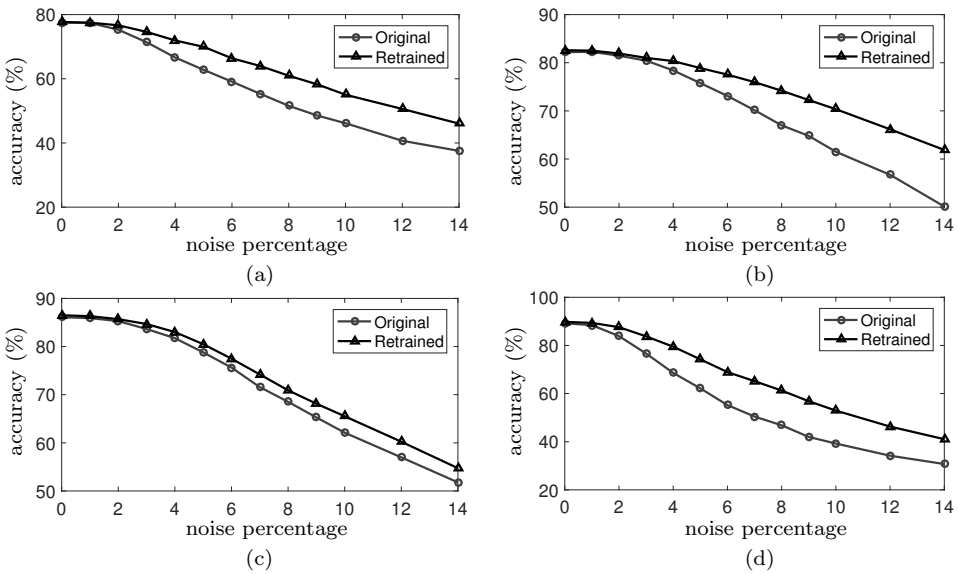

Figure 6: Classification accuracy for various noise levels before and after Net-Trim retraining: (a) 200 training samples; (b) 300 training samples; (c) 500 training samples; (d) 1000 training samples

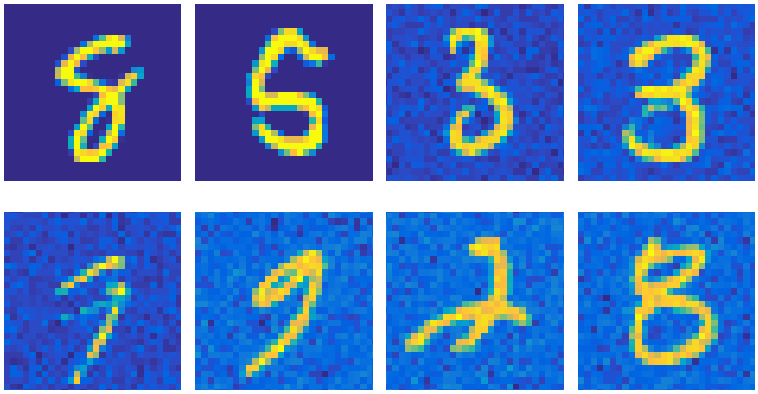

Figure 7: Examples of digits at different noise levels, which are misidentified using the initial trained network (with 1000 samples), but correctly identified after applying Net-Trim; first and second row initial identifications, left to right are 5, 6, 8, 8, 7, 7, 3, 6; the correct identified digits after applying the Net-Trim are 8, 5, 3, 3, 9, 9, 2, 8

6 Discussion and Conclusion

In linear regression problems, a well-known ℓ_1 regularization tool to promote sparsity on the resulting weights is the LASSO (least absolute shrinkage and selection operator [21]). The convexity of the initial regression problem allows us to conveniently add the ℓ_1 penalty to the problem. The convex structure of the ℓ_1-constrained problem helps relating the number of training samples to the active features in the model, as mainly presented in the compressed sensing literature [12]. In the neural networks, such process is not easy as the initial learning problem is highly non-convex. In this case, not much could be said in terms of characterizing the minimizers and the quality of the resulting models after adding an ℓ_1 regularizer.

By taking a layer-wise modeling strategy in Net-Trim, we have been able to overcome the non-convexity issue and present our results in an algorithmic and analytical way. The post-processing nature of the problem allows the algorithm to be conveniently blended with the state-of-the-art learning techniques in neural networks. Basically, regardless of the process being taken to train the model, Net-Trim can be considered as an additional post-processing step to reduce the model order and further improve the stability and prediction accuracy.

Similar to the LASSO, the proposed framework can handle both over and under-determined training cases and basically prevent networks of arbitrary size from overfitting. As expected, Net-Trim performs a more considerable pruning job on redundant networks, where the model degrees of freedom are more than those imposed by the data.

The Net-Trim pruning technique is not only limited to the parallel and cascade frameworks presented in this paper. As the method follows a layer-wise retraining strategy, it can be applied partially to a selected number of layers in the network. Even a hybrid retraining framework applying the parallel and cascade schemes to different subsets of the layers could be generally considered. To determine the optimal free parameters in the retrained model, cross validation methods can be used once the retraining scheme is set up.

From an optimization point of view, Net-Trim includes a series of constrained quadratic programs, which can be handled via any standard convex solver. A desirable characteristic of the resulting programs is the availability of at least one feasible point, which can be employed as an initialization for some solvers or minimization methods.

While the focus of this paper was solely an ℓ_1 regularization, with some slight modification to the formulation, other regularizations such as the ℓ_2 and max norm, or a combination of convex penalties such as $\ell_1 + \ell_2$ may be considered. In the majority of the cases, a slight change to the objective (e.g., replacing $\|U\|_1$ with $\|U\|_F$ or $\|U\|_1 + \lambda\|U\|_F$) brings the desired structure to the retrained models.

In the case of big data, where the network size and the training samples are large, Net-Trim presents a good compatibility. As stated in the paper, not only the parallel scheme allows retraining the layers individually and through independent computational resources, but also within a single layer retraining framework we can consider a PCN approach to cast the problem as a series of smaller independent programs. Yet, in cases that the number of

training samples are huge, an additional reduction in computational load would be to use the validation set instead of the training set in the Net-Trim. In a standard training phase, the update on the network weights stops once the prediction error within a reference validation set starts to increase. This set is often much smaller than the original training set and may be replaced with the training data in Net-Trim.

7 Appendix: Proofs of the Main Results

7.1 Proof of Proposition 1

If $\hat{\boldsymbol{W}} = [\hat{\boldsymbol{w}}_1, \cdots, \hat{\boldsymbol{w}}_M]$ is a solution to (6), then the feasibility of the solution requires

$$\sum_{m,p\,:\,y_{m,p}>0} (\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p - y_{m,p})^2 \le \epsilon^2 \tag{30}$$

and

$$\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p \le 0 \qquad \text{if} \quad y_{m,p} = 0. \tag{31}$$

Consider $\hat{\boldsymbol{Y}} = [\hat{y}_{m,p}]$, then

$$\begin{aligned}
\left\|\boldsymbol{Y} - \hat{\boldsymbol{Y}}\right\|_F^2 &= \sum_{m=1}^{M}\sum_{p=1}^{P}(y_{m,p} - \hat{y}_{m,p})^2 \\
&= \sum_{m,p\,:\,y_{m,p}>0} (y_{m,p} - \hat{y}_{m,p})^2 + \sum_{m,p\,:\,y_{m,p}=0} (y_{m,p} - \hat{y}_{m,p})^2 \\
&= \sum_{m,p\,:\,y_{m,p}>0} \left(y_{m,p} - (\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p)^+\right)^2 \\
&= \sum_{m,p\,:\,y_{m,p}>0,\hat{\boldsymbol{w}}_m^T\boldsymbol{x}_p>0} \left(y_{m,p} - \hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2 + \sum_{m,p\,:\,y_{m,p}>0,\hat{\boldsymbol{w}}_m^T\boldsymbol{x}_p\le 0} y_{m,p}^2.
\end{aligned} \tag{32}$$

Here since $y_{m,p} \ge 0$, the second equality is partitioned into two summations separated by the values of $y_{m,p}$ being zero or strictly greater than zero. The second resulting sum vanishes in the third equality since from (31), $\hat{\boldsymbol{y}}_{m,p} = \max(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p, 0) = 0$ when $\boldsymbol{y}_{m,p} = 0$. For the second term in (32) we use the basic algebraic identity

$$\sum_{m,p\,:\,y_{m,p}>0,\hat{\boldsymbol{w}}_m^T\boldsymbol{x}_p\le 0} y_{m,p}^2 = \sum_{m,p\,:\,y_{m,p}>0,\hat{\boldsymbol{w}}_m^T\boldsymbol{x}_p\le 0} \left(y_{m,p} - \hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2 + 2y_{m,p}(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p) - \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2. \tag{33}$$

Combining (33) and (32) results in

$$\left\|\boldsymbol{Y} - \hat{\boldsymbol{Y}}\right\|_F^2 = \sum_{m,p\,:\,y_{m,p}>0} \left(y_{m,p} - \hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2 + \sum_{m,p\,:\,y_{m,p}>0,\hat{\boldsymbol{w}}_m^T\boldsymbol{x}_p\le 0} 2y_{m,p}(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p) - \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2. \tag{34}$$

From (30), the first sum in (34) is upper bounded by ϵ^2. In addition,

$$2y_{m,p}(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p) - \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p\right)^2 \le 0,$$

when $y_{m,p} > 0$ and $\hat{\boldsymbol{w}}_m^T \boldsymbol{x}_p \le 0$, which together yield $\|\boldsymbol{Y} - \hat{\boldsymbol{Y}}\|_F^2 \le \epsilon^2$ as expected.

7.2 Proof of Theorem 1

We prove the theorem by induction. For $\ell = 1$, the claim holds as a direct result of Proposition 1. Now suppose the claim holds up to the $(\ell-1)$-th layer,

$$\left\|\hat{\boldsymbol{Y}}^{(\ell-1)} - \boldsymbol{Y}^{(\ell-1)}\right\|_F \le \sum_{j=1}^{\ell-1} \varepsilon_j, \tag{35}$$

we show that (9) will hold for the layer ℓ as well. The outcome of the ℓ-th layer before and after retraining obeys

$$y_{m,p}^{(\ell)} = \left(\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^+ \qquad \text{and} \qquad \hat{y}_{m,p}^{(\ell)} = \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+, \tag{36}$$

where $y_{m,p}^{(\ell)}$ and $\hat{y}_{m,p}^{(\ell)}$ are entries of $\boldsymbol{Y}^{(\ell)}$ and $\hat{\boldsymbol{Y}}^{(\ell)}$, the m-th columns of $\boldsymbol{W}_\ell$ and $\hat{\boldsymbol{W}}_\ell$ are denoted by $\boldsymbol{w}_m$ and $\hat{\boldsymbol{w}}_m$ (we have dropped the ℓ subscripts in the column notation for brevity), and the p-th columns of $\boldsymbol{Y}^{(\ell-1)}$ and $\hat{\boldsymbol{Y}}^{(\ell-1)}$ are denoted by $\boldsymbol{y}_p^{(\ell-1)}$ and $\hat{\boldsymbol{y}}_p^{(\ell-1)}$. We also define the quantities

$$\tilde{y}_{m,p}^{(\ell)} = \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^+,$$

which form a matrix $\tilde{\boldsymbol{Y}}^{(\ell)}$. From Proposition 1, we have

$$\left\|\tilde{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\right\|_F \le \epsilon_\ell. \tag{37}$$

On the other hand,

$$\begin{aligned}
\hat{y}_{m,p}^{(\ell)} &= \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ \\
&= \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{y}_p^{(\ell-1)} + \hat{\boldsymbol{w}}_m^T \left(\hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{y}_p^{(\ell-1)}\right)\right)^+ \\
&\le \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^+ + \left(\hat{\boldsymbol{w}}_m^T \left(\hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{y}_p^{(\ell-1)}\right)\right)^+ \\
&\le \tilde{y}_{m,p}^{(\ell)} + \left|\hat{\boldsymbol{w}}_m^T \left(\hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{y}_p^{(\ell-1)}\right)\right|,
\end{aligned} \tag{38}$$

where in the last two inequalities we used the sub-additivity of the $\max(.,0)$ function and the inequality $\max(x,0) \le |x|$. In a similar fashion we have

$$\begin{aligned}\tilde{y}_{m,p}^{(\ell)} &= \left(\hat{\boldsymbol{w}}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^+ \\ &\le \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ + \left(\hat{\boldsymbol{w}}_m^T \left(\boldsymbol{y}_p^{(\ell-1)} - \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)\right)^+ \\ &\le \hat{y}_{m,p}^{(\ell)} + \left|\hat{\boldsymbol{w}}_m^T \left(\hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{y}_p^{(\ell-1)}\right)\right|,\end{aligned}$$

which together with (38) asserts that $|\hat{y}_{m,p}^{(\ell)} - \tilde{y}_{m,p}^{(\ell)}| \le |\hat{\boldsymbol{w}}_m^T(\hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{y}_p^{(\ell-1)})|$ or

$$\left\|\hat{\boldsymbol{Y}}^{(\ell)} - \tilde{\boldsymbol{Y}}^{(\ell)}\right\|_F \le \left\|\hat{\boldsymbol{W}}_\ell^T \left(\hat{\boldsymbol{Y}}^{(\ell-1)} - \boldsymbol{Y}^{(\ell-1)}\right)\right\|_F \le \left\|\hat{\boldsymbol{W}}_\ell\right\|_F \left\|\hat{\boldsymbol{Y}}^{(\ell-1)} - \boldsymbol{Y}^{(\ell-1)}\right\|_F. \tag{39}$$

As $\hat{\boldsymbol{W}}_\ell$ is the minimizer of (8) and $\boldsymbol{W}_\ell$ is a feasible point (i.e., $\boldsymbol{W}_\ell \in \mathcal{C}_{\epsilon_\ell}(\boldsymbol{Y}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{0})$), we have

$$\|\hat{\boldsymbol{W}}_\ell\|_F \le \|\hat{\boldsymbol{W}}_\ell\|_1 \le \|\boldsymbol{W}_\ell\|_1 = 1, \tag{40}$$

which with reference to (39) yields

$$\left\|\hat{\boldsymbol{Y}}^{(\ell)} - \tilde{\boldsymbol{Y}}^{(\ell)}\right\|_F \le \left\|\hat{\boldsymbol{Y}}^{(\ell-1)} - \boldsymbol{Y}^{(\ell-1)}\right\|_F \le \sum_{j=1}^{\ell-1} \varepsilon_j.$$

Finally, the induction proof is completed by applying the triangle inequality and then using (37),

$$\left\|\hat{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\right\|_F \le \left\|\hat{\boldsymbol{Y}}^{(\ell)} - \tilde{\boldsymbol{Y}}^{(\ell)}\right\|_F + \left\|\tilde{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\right\|_F \le \sum_{j=1}^{\ell} \varepsilon_j.$$

7.3 Proof of Theorem 2

For $\ell \ge 2$ we relate the upper-bound of $\|\hat{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\|_F$ to $\|\hat{\boldsymbol{Y}}^{(\ell-1)} - \boldsymbol{Y}^{(\ell-1)}\|_F$. By the construction of the network:

$$\begin{aligned}\left\|\hat{\boldsymbol{Y}}^{(\ell)} - \boldsymbol{Y}^{(\ell)}\right\|_F^2 &= \sum_{m=1}^{N_\ell} \sum_{p=1}^{P} \left(\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ - \left(\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^+\right)^2 \\ &= \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0} \left(\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 + \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0} \left(\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+\right)^2, \end{aligned} \tag{41}$$

where the m-th columns of $\boldsymbol{W}_\ell$ and $\hat{\boldsymbol{W}}_\ell$ are denoted by $\boldsymbol{w}_m$ and $\hat{\boldsymbol{w}}_m$, respectively (we have dropped the ℓ subscripts in the column notation for brevity), and the p-th columns of $\boldsymbol{Y}^{(\ell)}$

and $\hat{\boldsymbol{Y}}^{(\ell)}$ are denoted by $\boldsymbol{y}_p^{(\ell-1)}$ and $\hat{\boldsymbol{y}}_p^{(\ell-1)}$. Also, as before $y_{m,p}^{(\ell)} = (\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)})^+$. For the first term in (41) we have

$$\sum_{m,p\,:\,y_{m,p}^{(\ell)}>0} \left(\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 = \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0, \hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}>0} \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 + \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0,\, \hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\le 0} \left(\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2. \tag{42}$$

However, for the elements of the set $\{(m,p) : y_{m,p}^{(\ell)} > 0,\ \hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} \le 0\}$:

$$\begin{aligned}\left(\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 &= \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 + 2\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)\left(\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right) - \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^2 \\ &\le \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2,\end{aligned}$$

using which in (42) yields

$$\begin{aligned}\sum_{m,p\,:\,y_{m,p}^{(\ell)}>0} \left(\left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+ - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 &\le \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0} \left(\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 \\ &\le \gamma_\ell \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0} \left(\boldsymbol{w}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} - \boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)}\right)^2 \\ &= \epsilon_\ell^2.\end{aligned} \tag{43}$$

Here, the second inequality is a direct result of the feasibility condition

$$\hat{\boldsymbol{W}}_\ell \in \mathcal{C}_{\epsilon_\ell}\left(\hat{\boldsymbol{Y}}^{(\ell-1)}, \boldsymbol{Y}^{(\ell)}, \boldsymbol{W}_\ell \hat{\boldsymbol{Y}}^{(\ell-1)}\right). \tag{44}$$

A second outcome of the feasibility is

$$\hat{\boldsymbol{w}}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)} \le \boldsymbol{w}_m^T \hat{\boldsymbol{y}}_p^{(\ell-1)}, \tag{45}$$

for any pair (m,p) that obeys $y_{m,p}^{(\ell)} = 0$ (or equivalently, $\boldsymbol{w}_m^T \boldsymbol{y}_p^{(\ell-1)} \le 0$). We can apply $\max(.,0)$ (as an increasing and positive function) to both sides of (45) and use it to bound the second

term in (41) as follows:

$$\begin{aligned}
\sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\left(\hat{\boldsymbol{w}}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+\right)^2 &\le \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}\right)^+\right)^2\\
&= \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\left(\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}+\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^+\right)^2\\
&\le \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\left(\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^+ + \left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^+\right)^2\\
&= \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^+\right)^2\\
&\le \sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^2. \qquad (46)
\end{aligned}$$

The first and second terms in (41) are bounded via (43) and (46) and therefore

$$\begin{aligned}
\left\|\hat{\boldsymbol{Y}}^{(\ell)}-\boldsymbol{Y}^{(\ell)}\right\|_F^2 &\le \gamma_\ell \sum_{m,p\,:\,y_{m,p}^{(\ell)}>0}\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^2+\sum_{m,p\,:\,y_{m,p}^{(\ell)}=0}\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^2\\
&\le \gamma_\ell \sum_{m,p}\left(\boldsymbol{w}_m^T\hat{\boldsymbol{y}}_p^{(\ell-1)}-\boldsymbol{w}_m^T\boldsymbol{y}_p^{(\ell-1)}\right)^2\\
&= \gamma_\ell \left\|\boldsymbol{W}_\ell^T\left(\hat{\boldsymbol{Y}}^{(\ell-1)}-\boldsymbol{Y}^{(\ell-1)}\right)\right\|_F^2\\
&\le \gamma_\ell \left\|\boldsymbol{W}_\ell\right\|_F^2\left\|\hat{\boldsymbol{Y}}^{(\ell-1)}-\boldsymbol{Y}^{(\ell-1)}\right\|_F^2\\
&= \gamma_\ell \left\|\hat{\boldsymbol{Y}}^{(\ell-1)}-\boldsymbol{Y}^{(\ell-1)}\right\|_F^2. \qquad (47)
\end{aligned}$$

Based on Proposition 1, the outcome of the first layer obeys $\|\hat{\boldsymbol{Y}}^{(1)}-\boldsymbol{Y}^{(1)}\|_F^2 \le \epsilon_1^2$, which together with (47) confirm (16).

7.4 Proof of Proposition 2

For part (a), from Theorem 1 we have $\|\hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L-1)}\|_F \le \sum_{\ell=1}^{L-1} \epsilon_\ell$. Furthermore,

$$\begin{aligned}
\left\|\hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)}\right\|_F &= \left\|\hat{\boldsymbol{W}}_L^T \hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L)}\right\|_F \\
&\le \left\|\hat{\boldsymbol{W}}_L^T \hat{\boldsymbol{Y}}^{(L-1)} - \hat{\boldsymbol{W}}_L^T \boldsymbol{Y}^{(L-1)}\right\|_F + \left\|\hat{\boldsymbol{W}}_L^T \boldsymbol{Y}^{(L-1)} - \boldsymbol{Y}^{(L)}\right\|_F \\
&\le \left\|\hat{\boldsymbol{W}}_L^T\right\|_F \left\|\hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L-1)}\right\|_F + \epsilon_L \\
&\le \sum_{\ell=1}^{L} \epsilon_\ell,
\end{aligned}$$

where for the last inequality we used a similar chain of inequalities as (40).

To prove part (b), for the first $L-1$ layers, $\|\hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L-1)}\|_F \le \epsilon_1 \sqrt{\prod_{\ell=1}^{L-1} \gamma_\ell}$, and therefore

$$\begin{aligned}
\left\|\hat{\boldsymbol{Y}}^{(L)} - \boldsymbol{Y}^{(L)}\right\|_F &= \left\|\hat{\boldsymbol{W}}_L^T \hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L)}\right\|_F \\
&\le \sqrt{\gamma_L} \left\|\boldsymbol{W}_L^T \hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{W}_L^T \boldsymbol{Y}^{(L-1)}\right\|_F \\
&\le \sqrt{\gamma_L}\, \|\boldsymbol{W}_L\|_F \left\|\hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L-1)}\right\|_F \\
&\le \sqrt{\gamma_L} \left\|\hat{\boldsymbol{Y}}^{(L-1)} - \boldsymbol{Y}^{(L-1)}\right\|_F \\
&\le \epsilon_1 \sqrt{\prod_{\ell=1}^{L} \gamma_\ell}.
\end{aligned}$$

Here, the first inequality is thanks to (18) and (19).

7.5 Proof of Proposition 3

Since program (21) has only affine inequality constraints, the Karush–Kuhn–Tucker (KKT) optimality conditions give us necessary and sufficient conditions for an optimal solution. The pair $(\boldsymbol{w}^*, \boldsymbol{s}^*)$ is optimal if and only if there exists $\boldsymbol{\eta} \in \mathbb{R}^{n_2}$ and $\boldsymbol{\nu}$ such that

$$\begin{aligned}
\eta_k s_k^* &= 0, \quad k = 1, \dots, n_2, \\
\eta_k &\ge 0, \quad k = 1, \dots, n_2, \\
\tilde{\boldsymbol{X}}^T \boldsymbol{\nu} &\in \begin{bmatrix} \partial \|\boldsymbol{w}^*\|_1 \\ \boldsymbol{\eta} \end{bmatrix}.
\end{aligned}$$

Above, $\partial\|\boldsymbol{w}^*\|_1$ denotes the subgradient of the ℓ_1 norm evaluated at $\boldsymbol{w}^*$; the last expression above means that the first n_1 entries of $\tilde{\boldsymbol{X}}^T\boldsymbol{\nu}$ must match the sign of w_ℓ^* for indexes ℓ with $w_\ell^* \neq 0$, and must have magnitude no greater than 1 for indexes ℓ with $w_\ell^* = 0$. The existence of such $(\boldsymbol{\eta}, \boldsymbol{\nu})$ is compatible with the existence of a $\boldsymbol{\Lambda}$ meeting the conditions in (22), by taking $\boldsymbol{\Lambda} = \tilde{\boldsymbol{X}}^T\boldsymbol{\nu}$.

We now argue the conditions for uniqueness. Let $\boldsymbol{w}^*, \boldsymbol{s}^*, \boldsymbol{\Lambda}$ be as above. Suppose $(\boldsymbol{w}', \boldsymbol{s}')$ is a feasible point with $\|\boldsymbol{w}'\|_1 = \|\boldsymbol{w}^*\|_1$. Since $\boldsymbol{\Lambda}$ is in the row space of $\tilde{\boldsymbol{X}}$, we know that

$$\boldsymbol{\Lambda}^T \begin{bmatrix} \boldsymbol{w}^* - \boldsymbol{w}' \\ \boldsymbol{s}^* - \boldsymbol{s}' \end{bmatrix} = 0,$$

and since $\boldsymbol{\Lambda}^T \left[\boldsymbol{w}^*; \boldsymbol{s}^*\right] = \|\boldsymbol{w}^*\|_1$, we must also have $\boldsymbol{\Lambda}^T \left[\boldsymbol{w}'; \boldsymbol{s}'\right] = \|\boldsymbol{w}^*\|_1$. Therefore by the properties stated in (22), the support (set of nonzero entries) $\tilde{\Gamma}$ of $[\boldsymbol{w}^*; \boldsymbol{s}^*]$ and $[\boldsymbol{w}'; \boldsymbol{s}']$ must be the same. Since these points obey the same equality constraints in the program (21), and $\tilde{\boldsymbol{X}}_{:,\tilde{\Gamma}}$ has full column rank, it must be true that $[\boldsymbol{w}'; \boldsymbol{s}'] = [\boldsymbol{w}^*; \boldsymbol{s}^*]$.

The interested reader is referred to [1, 2] for examples of more comprehensive uniqueness proofs for convex problems with both equality and inequality constraints.

7.6 Proof of Theorem 3

For a more convenient notation we use $\Gamma = \operatorname{supp} \boldsymbol{w}^*$. Also, in all the formulations, sub-matrix selection precedes the transpose operation, e.g., $\boldsymbol{X}_{:,\Omega}^T = (\boldsymbol{X}_{:,\Omega})^T$.

Clearly since the samples are random Gaussians, with probability one the set $\{p : \boldsymbol{X}_{:,p}^T \boldsymbol{w}^* = 0\}$ is an empty set, and following the notation in (21) and (22), $\operatorname{supp}^c \boldsymbol{s}^* = \varnothing$. Also, with reference to the setup in Proposition 3

$$\tilde{\boldsymbol{X}} = \begin{bmatrix} \boldsymbol{X}_{:,\Omega}^T & \boldsymbol{0} \\ \boldsymbol{X}_{:,\Omega^c}^T & -\boldsymbol{I} \end{bmatrix}.$$

To establish the full column rank property of $\tilde{\boldsymbol{X}}_{:,\tilde{\Gamma}}$ for $\tilde{\Gamma} = \operatorname{supp} \boldsymbol{w}^* \cup \{N + \operatorname{supp} \boldsymbol{s}^*\}$, we only need to show that $\boldsymbol{X}_{\Gamma,\Omega}$ is of full row rank (thanks to the identity block in $\tilde{\boldsymbol{X}}$). Also, to satisfy the equality requirements in (22), we need to find a vector $\boldsymbol{\xi}$ such that

$$\begin{bmatrix} \boldsymbol{X}_{\Gamma,\Omega} & \boldsymbol{X}_{\Gamma,\Omega^c} \\ \boldsymbol{0} & -\boldsymbol{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi}_\Omega \\ \boldsymbol{\xi}_{\Omega^c} \end{bmatrix} = \begin{bmatrix} \operatorname{sign}(\boldsymbol{w}_\Gamma^*) \\ \boldsymbol{0} \end{bmatrix}. \tag{48}$$

This equation trivially yields $\boldsymbol{\xi}_{\Omega^c} = \boldsymbol{0}$. In the remainder of the proof we will show that when P is sufficiently larger than $|\Gamma| = s$, the smallest eigenvalue of $\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T$ is bounded away from zero (which automatically establishes the full row rank property for $\boldsymbol{X}_{\Gamma,\Omega}$). Also, based on such property, we can select $\boldsymbol{\xi}_\Omega$ to be the least squares solution

$$\boldsymbol{\xi}_\Omega \triangleq \boldsymbol{X}_{\Gamma,\Omega}^T \left(\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\right)^{-1} \operatorname{sign}(\boldsymbol{w}_\Gamma^*), \tag{49}$$

which satisfies the equality condition in (48). To verify the conditions stated in (22) and complete the proof, we will show that when P is sufficiently large, with high probability $\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty < 1$. We do this by bounding the probability of failure via the inequality

$$\mathbb{P}\left(\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1\right) \leq \mathbb{P}\left(\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1 \;\middle|\; \|\boldsymbol{\xi}_\Omega\| \leq \tau\right) + \mathbb{P}\left(\|\boldsymbol{\xi}_\Omega\| > \tau\right), \tag{50}$$

for some positive τ, which is a simple consequence of

$$\mathbb{P}\left(\mathcal{E}_1\right) = \mathbb{P}\left(\mathcal{E}_1|\mathcal{E}_2\right)\mathbb{P}\left(\mathcal{E}_2\right) + \mathbb{P}\left(\mathcal{E}_1|\mathcal{E}_2^c\right)\mathbb{P}\left(\mathcal{E}_2^c\right) \leq \mathbb{P}\left(\mathcal{E}_1|\mathcal{E}_2\right) + \mathbb{P}\left(\mathcal{E}_2^c\right),$$

generally holding for two event $\mathcal{E}_1$ and $\mathcal{E}_2$. Without the filtering of the Ω set, standard concentration bounds on the least squares solution can help establishing the unique optimality conditions (e.g., see [9]). Here also, we proceed by bounding each term on the right hand side of (50) individually, while the bounding process requires taking a different path because of the dependence Ω brings to the resulting random matrices.

- **Step 1. Bounding $\mathbb{P}(\|\boldsymbol{\xi}_\Omega\| > \tau)$:**

By the construction of $\boldsymbol{\xi}_\Omega$ in(49), clearly

$$\|\boldsymbol{\xi}_\Omega\|^2 = \operatorname{sign}(\boldsymbol{w}_\Gamma^*)^T\left(\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\right)^{-1}\operatorname{sign}(\boldsymbol{w}_\Gamma^*). \tag{51}$$

Technically speaking, to bound the expression $\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x}$, where $\boldsymbol{x}$ is a fixed vector and $\boldsymbol{A}$ is a self adjoint random matrix, we normally need the entries of $\boldsymbol{A}$ to be independent of the elements in $\boldsymbol{x}$. While such independence does not hold in (51) (because of the dependence of Ω to the entries of $\boldsymbol{w}_\Gamma^*$), we are still able to proceed with bounding by rewriting $\operatorname{sign}(\boldsymbol{w}_\Gamma^*) = \boldsymbol{\Lambda}_{\boldsymbol{w}^*}\mathbf{1}$, where

$$\boldsymbol{\Lambda}_{\boldsymbol{w}^*} = \operatorname{diag}\left(\operatorname{sign}(\boldsymbol{w}_\Gamma^*)\right).$$

Taking into account the facts that $\boldsymbol{\Lambda}_{\boldsymbol{w}^*} = \boldsymbol{\Lambda}_{\boldsymbol{w}^*}^{-1}$ and $w_n^* \neq 0$ for $n \in \Gamma$, we have

$$\|\boldsymbol{\xi}_\Omega\|^2 = \mathbf{1}^T\left(\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\right)^{-1}\mathbf{1}, \tag{52}$$

where now the matrix and vector independence is maintained. The special structure of $\boldsymbol{\Lambda}_{\boldsymbol{w}^*}$ does not cause a change in the eigenvalues and

$$\operatorname{eig}\left\{\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\right\} = \operatorname{eig}\left\{\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\right\},$$

where eig{.} denotes the set of eigenvalues. Now conditioned on $\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T \succ \mathbf{0}$, we can bound the magnitude of $\boldsymbol{\xi}_\Omega$ as

$$\begin{aligned}
\|\boldsymbol{\xi}_\Omega\|^2 &= \mathbf{1}^T\left(\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\right)^{-1}\mathbf{1} \\
&\leq \lambda_{\max}\left(\left(\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\boldsymbol{\Lambda}_{\boldsymbol{w}^*}\right)^{-1}\right)\mathbf{1}^T\mathbf{1} \\
&= s\left(\lambda_{\min}\left(\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\right)\right)^{-1},
\end{aligned} \tag{53}$$

where $\lambda_{\max}$ and $\lambda_{\min}$ denote the maximum and minimum eigenvalues. To lower bound $\lambda_{\min}\left(\boldsymbol{X}_{\Gamma,\Omega}\boldsymbol{X}_{\Gamma,\Omega}^T\right)$, we focus on the matrix eigenvalue results associated with the sum of random matrices. For this purpose, consider the independent sequence of random vectors $\{\boldsymbol{x}_p\}_{p=1}^P$, where each vector contains i.i.d standard normal entries. We are basically interested in concentration bounds for

$$\lambda_{\min}\left(\sum_{p\,:\,\boldsymbol{x}_p^T\boldsymbol{w}^*>0}\boldsymbol{x}_p\boldsymbol{x}_p^T\right). \tag{54}$$

When the summands are independent self adjoint random matrices, we can use standard Bernstein type inequalities to bound the minimum or maximum eigenvalues [22]. However, as the summands in (54) are dependent in the sense that they all obey $\boldsymbol{x}_p^T\boldsymbol{w}^* > 0$, such results are not directly applicable. To establish the independence, we can look into an equivalent formulation of (54) as

$$\lambda_{\min}\left(\sum_{p=1}^{P}\boldsymbol{x}_p\boldsymbol{x}_p^T\right), \tag{55}$$

where $\boldsymbol{x}_p$ are independently drawn from the distribution

$$g_{\boldsymbol{X}}(\boldsymbol{x}) = \begin{cases} \frac{1}{\sqrt{(2\pi)^s}}\exp\left(-\frac{1}{2}\boldsymbol{x}^T\boldsymbol{x}\right) + \frac{1}{2}\delta_D(\boldsymbol{x}) & \boldsymbol{x}^T\boldsymbol{w}^* \geq 0 \\ 0 & \boldsymbol{x}^T\boldsymbol{w}^* < 0 \end{cases}.$$

Here, $\delta_D(\boldsymbol{x}) = \prod_{i=1}^s \delta_D(x_i)$ denotes the s-dimensional Dirac delta function, and is probabilistically in charge of returning a zero vector in half of the draws. We are now theoretically able to apply the following result, brought from [22], to bound the smallest eigenvalue:

Theorem 4. *(Matrix Bernstein[2]) Consider a finite sequence $\boldsymbol{Z}_p$ of independent, random, self-adjoint matrices with dimension s. Assume that each random matrix satisfies*

$$\mathbb{E}(\boldsymbol{Z}_p) = \boldsymbol{0}, \quad \textit{and} \quad \lambda_{\min}(\boldsymbol{Z}_p) \geq R \quad \textit{almost surely.}$$

Then, for all $t \leq 0$,

$$\mathbb{P}\left\{\lambda_{\min}\left(\sum_p \boldsymbol{Z}_p\right) \leq t\right\} \leq s\exp\left(\frac{-t^2}{2\sigma^2 + 2Rt/3}\right),$$

where $\sigma^2 = \|\sum_p \mathbb{E}(\boldsymbol{Z}_p^2)\|$.

To more conveniently apply Theorem 4, we can use a change of variable which markedly simplifies the moment calculations required for the Bernstein inequality. For this purpose,

[2]The original version of the theorem bounds the maximum eigenvalue. The present version can be easily derived using, $\lambda_{\min}(\boldsymbol{Z}) = -\lambda_{\max}(-\boldsymbol{Z})$ and $\mathbb{P}\{\lambda_{\min}(\sum_p \boldsymbol{Z}_p) \leq t\} = \mathbb{P}\{\lambda_{\max}(\sum_p -\boldsymbol{Z}_p) \geq -t\}$.

consider $\boldsymbol{R}$ to be a rotation matrix which maps $\boldsymbol{w}^*$ to the first canonical basis $[1, 0, \cdots, 0]^T \in \mathbb{R}^s$. Since

$$\text{eig}\left\{\sum_{p=1}^{P} \boldsymbol{x}_p \boldsymbol{x}_p^T\right\} = \text{eig}\left\{\sum_{p=1}^{P} \boldsymbol{R}\boldsymbol{x}_p \boldsymbol{x}_p^T \boldsymbol{R}^T\right\}, \tag{56}$$

we can focus on random vectors $\boldsymbol{u}_p = \boldsymbol{R}\boldsymbol{x}_p$ which follow the simpler distribution

$$g_{\boldsymbol{U}}(\boldsymbol{u}) \triangleq \begin{cases} \frac{1}{\sqrt{(2\pi)^s}} \exp\left(-\frac{1}{2}\boldsymbol{u}^T\boldsymbol{u}\right) + \frac{1}{2}\delta_D(\boldsymbol{u}) & u_1 \geq 0 \\ 0 & u_1 < 0 \end{cases}. \tag{57}$$

Here, u_1 denotes the first entry of $\boldsymbol{u}$, and we used the basic property $\boldsymbol{R}^{-1} = \boldsymbol{R}^T$ along with the rotation invariance of the Dirac delta function to derive $g_{\boldsymbol{U}}(\boldsymbol{u})$ from $g_{\boldsymbol{X}}(\boldsymbol{x})$. Using the Bernstein inequality, we can now summarize everything as the following concentration result (proved later in the section):

Proposition 4. *Consider a sequence of independent $\{\boldsymbol{u}_p\}_{p=1}^P$ vectors of length s, where each vector is drawn from the distribution $g_{\boldsymbol{U}}(\boldsymbol{u})$ in (57). For all $t \leq 0$,*

$$\mathbb{P}\left\{\lambda_{\min}\left(\sum_{p=1}^{P} \boldsymbol{u}_p \boldsymbol{u}_p^T\right) \leq P + t\right\} \leq s \exp\left(\frac{-t^2}{P(2s+1) - 2t/3}\right). \tag{58}$$

Combining the lower bound in (53) with the concentration result (58) certify that when $P + t > 0$ and $t \leq 0$,

$$\mathbb{P}\left\{\|\boldsymbol{\xi}_\Omega\| \geq \sqrt{\frac{s}{P+t}}\right\} \leq s \exp\left(\frac{-t^2}{P(2s+1) - 2t/3}\right). \tag{59}$$

- **Step 2. Bounding $\mathbb{P}\{\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1 \mid \|\boldsymbol{\xi}_\Omega\| \leq \tau\}$:**

Considering the conditioned event $\{\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1 \mid \|\boldsymbol{\xi}_\Omega\| \leq \tau\}$, we note that the set Ω is constructed by selecting columns of $\boldsymbol{X}$ that satisfy $\boldsymbol{X}_{:,p}^T \boldsymbol{w}^* > 0$. However, since $\boldsymbol{w}^*_{\Gamma^c} = \boldsymbol{0}$, the index set Ω, technically corresponds to the columns p where $\boldsymbol{X}_{\Gamma,p}^T \boldsymbol{w}^*_\Gamma > 0$. In other words, none of the entries of the sub-matrix $\boldsymbol{X}_{\Gamma^c,:}$ contribute to the selection of Ω. Noting this, conditioned on given $\boldsymbol{\xi}_\Omega$, the entries of the vector $\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega$ are i.i.d random variables distributed as $\mathcal{N}(0, \|\boldsymbol{\xi}_\Omega\|^2)$ and

$$\mathbb{P}\left\{\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1 \mid \|\boldsymbol{\xi}_\Omega\| \leq \tau\right\} = \mathbb{P}\left\{\bigcup_{n=1}^{|\Gamma^c|} |z_n| \geq \frac{1}{\|\boldsymbol{\xi}_\Omega\|} \;\middle|\; \|\boldsymbol{\xi}_\Omega\| \leq \tau\right\}, \tag{60}$$

where $\{z_n\}$ are i.i.d standard normals. Using the union bound and the basic inequality $\mathbb{P}(|z_n| \geq a) \leq \exp(-a^2/2)$ valid for $a \geq 0$, we get

$$\mathbb{P}\left\{\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1 \mid \|\boldsymbol{\xi}_\Omega\| \leq \tau\right\} \leq (N - s) \exp\left(-\frac{1}{2\tau^2}\right). \tag{61}$$

For $\tau = \sqrt{s(P+t)^{-1}}$ we can combine (61) and (59) with reference to (50) to get

$$\mathbb{P}\left(\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1\right) \leq s\exp\left(\frac{-t^2}{P(2s+1)-2t/3}\right) + (N-s)\exp\left(-\frac{P+t}{2s}\right). \tag{62}$$

To select the free parameter t we make the argument of the two exponentials equal to get

$$t^* = \frac{6s+1-\sqrt{180s^2+132s+25}}{12s+4}P,$$

for which the right hand side expression in (62) reduces to $N\exp(-(2s)^{-1}(P+t^*))$. Based on the given value of t^*, it is easy to verify that for all $P \geq 0$ and $s \geq 1$, the conditions $t^* \leq 0$ and $P + t^* > 0$ are satisfied. Moreover some basic algebra reveals that for all $P \geq 0$ and $s \geq 1$

$$-\frac{P+t^*}{2s} \leq -\frac{P}{6s+1}.$$

Therefore, for $\mu > 1$, setting $P = (6s+1)\mu\log N$ guarantees that

$$\mathbb{P}\left(\|\boldsymbol{X}_{\Gamma^c,\Omega}\boldsymbol{\xi}_\Omega\|_\infty \geq 1\right) \leq N^{\mu-1}.$$

7.6.1 Proof of Proposition 4

To use Theorem 4, we focus on a sequence $\{\boldsymbol{Z}_p\}_{p=1}^P$ of the random matrices $\boldsymbol{Z} = \boldsymbol{u}\boldsymbol{u}^T - \mathbb{E}(\boldsymbol{u}\boldsymbol{u}^T)$. In all the steps discussed below, we need to calculate cross moments of the type $\mathbb{E}_g(u_1^{n_1}u_2^{n_2}\cdots u_s^{n_s})$ for $\boldsymbol{u} = [u_i] \in \mathbb{R}^s$ distributed as

$$g_U(\boldsymbol{u}) \triangleq \begin{cases} \frac{1}{\sqrt{(2\pi)^s}}\exp\left(-\frac{1}{2}\boldsymbol{u}^T\boldsymbol{u}\right) + \frac{1}{2}\delta_D(\boldsymbol{u}) & u_1 \geq 0 \\ 0 & u_1 < 0 \end{cases}.$$

Instead we can calculate $\mathbb{E}(u_1^{n_1}u_2^{n_2}\cdots u_s^{n_s})$ of a random vector $\boldsymbol{u} = [u_i] \in \mathbb{R}^s$ with *independent* entries such that u_1 is distributed as

$$f_U(u) \triangleq \begin{cases} \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{u^2}{2}\right) + \frac{1}{2}\delta_D(u) & u \geq 0 \\ 0 & u < 0 \end{cases}. \tag{63}$$

and $u_i \sim \mathcal{N}(0,1)$ for $i = 2,\cdots,s$. In either case we have

$$\mathbb{E}_g(\prod_{i=1}^s u_i^{n_i}) = \mathbb{E}(\prod_{i=1}^s u_i^{n_i}) = \frac{1}{\sqrt{(2\pi)^s}}\int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty}\int_0^{\infty}\exp\left(-\frac{1}{2}\sum_{i=1}^s u_i^2\right)\prod_{i=1}^s u_i^{n_i}\mathrm{d}u_1\mathrm{d}u_2\cdots\mathrm{d}u_s,$$

simply because

$$\int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty}\delta_D(u_1)\prod_{i=1}^{s}u_i^{n_i}\mathrm{d}u_1\mathrm{d}u_2\cdots\mathrm{d}u_s=\int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty}\prod_{i=1}^{s}u_i^{n_i}\delta_D(u_i)\mathrm{d}u_1\mathrm{d}u_2\cdots\mathrm{d}u_s=0.$$

For u_1 distributed as $f_U(u)$, elementary calculation reveals that

$$\mathbb{E}\left(u_1^2\right)=\frac{1}{2},\qquad \mathbb{E}\left(u_1^4\right)=\frac{3}{2}. \tag{64}$$

We now partition $\boldsymbol{u}$ as

$$\boldsymbol{u}=\begin{bmatrix}u_1\\ \boldsymbol{u}_G\end{bmatrix},$$

where $\boldsymbol{u}_G=(u_2,\cdots,u_s)^T$ denotes the i.i.d normally distributed components. Using (64) and the independence of the components we get

$$\boldsymbol{D}=\mathbb{E}(\boldsymbol{u}\boldsymbol{u}^T)=\mathrm{diag}(\frac{1}{2},1,\cdots,1).$$

The matrix $\boldsymbol{u}\boldsymbol{u}^T$ is a rank one positive semidefinite matrix, which has only one nonzero eigenvalue. Using the Weyl's inequality we get

$$\lambda_{\min}\left(\boldsymbol{Z}\right)=\lambda_{\min}\left(\boldsymbol{u}\boldsymbol{u}^T-\boldsymbol{D}\right)\geq\lambda_{\min}(\boldsymbol{u}\boldsymbol{u}^T)+\lambda_{\min}(-\boldsymbol{D})=-1. \tag{65}$$

Furthermore,

$$\mathbb{E}\left(\boldsymbol{Z}_p^2\right)=\mathbb{E}\left(\left(\boldsymbol{u}\boldsymbol{u}^T\right)^2\right)-\boldsymbol{D}^2,$$

for which we can calculate the expectation term as

$$\begin{aligned}\mathbb{E}\left(\left(\boldsymbol{u}\boldsymbol{u}^T\right)^2\right)&=\begin{bmatrix}\mathbb{E}\left(u_1^4+u_1^2\boldsymbol{u}_G^T\boldsymbol{u}_G\right) & \mathbb{E}\left(u_1^3\boldsymbol{u}_G^T+u_1\boldsymbol{u}_G^T\boldsymbol{u}_G\boldsymbol{u}_G^T\right)\\ \mathbb{E}\left(u_1^3\boldsymbol{u}_G+u_1\boldsymbol{u}_G\boldsymbol{u}_G^T\boldsymbol{u}_G\right) & \mathbb{E}\left(u_1^2\boldsymbol{u}_G\boldsymbol{u}_G^T+(\boldsymbol{u}_G\boldsymbol{u}_G^T)^2\right)\end{bmatrix}\\ &=\begin{bmatrix}\frac{3}{2}+\frac{1}{2}(s-1) & \boldsymbol{0}^T\\ \boldsymbol{0} & \frac{1}{2}\boldsymbol{I}+(s+1)\boldsymbol{I}\end{bmatrix}.\end{aligned}$$

Here, we used (64) and the result of the following simple lemma:

Lemma 1. *Given a random vector $\boldsymbol{u}=[u_i]\in\mathbb{R}^s$, with i.i.d entries $u_i\sim\mathcal{N}(0,1)$:*

$$\mathbb{E}\left((\boldsymbol{u}\boldsymbol{u}^T)^2\right)=(s+2)\boldsymbol{I}.$$

It is now easy to observe that

$$\left\| \sum_{p=1}^{P} \mathbb{E}\left(\boldsymbol{Z}_p^2\right) \right\| = P\lambda_{\max}\left(\mathbb{E}\left(\left(\boldsymbol{u}\boldsymbol{u}^T\right)^2\right) - \boldsymbol{D}^2\right) = P\left(\frac{1}{2} + (s+1) - 1\right) = P\left(s + \frac{1}{2}\right). \tag{66}$$

Now, using (65) and (66) we can apply Theorem 4 to bound the smallest eigenvalue as

$$\mathbb{P}\left\{ \lambda_{\min}\left(\sum_{p=1}^{P} \boldsymbol{u}_p \boldsymbol{u}_p^T - P\boldsymbol{D} \right) \le t \right\} \le s \exp\left(\frac{-t^2}{P(2s+1) - 2t/3} \right), \tag{67}$$

for all $t \le 0$. The Weyl's inequality yields

$$\mathbb{P}\left\{ \lambda_{\min}\left(\sum_{p=1}^{P} \boldsymbol{u}_p \boldsymbol{u}_p^T \right) \le P + t \right\} \le \mathbb{P}\left\{ \lambda_{\min}\left(\sum_{p=1}^{P} \boldsymbol{u}_p \boldsymbol{u}_p^T - P\boldsymbol{D} \right) \le t \right\} \tag{68}$$

which gives the probability mentioned in (58).

7.6.2 Proof of Lemma 1

The (i,j) element of the underlying matrix can be written as

$$\left((\boldsymbol{u}\boldsymbol{u}^T)^2\right)_{i,j} = u_i u_j \sum_{k=1}^{s} u_k^2,$$

therefore,

$$\mathbb{E}\left(u_i u_j \sum_{k=1}^{s} u_k^2 \right) = \begin{cases} 0 & i \neq j \\ \mathbb{E}\left(u_i^4 + u_i^2 \sum_{k \neq i} u_k^2\right) & i = j \end{cases} = \begin{cases} 0 & i \neq j \\ s+2 & i = j \end{cases}. \tag{69}$$

Here, we used the facts that $\mathbb{E}(u_i^4) = 3$ and $\sum_{k \neq i} u_k^2 = s - 1$.

References

[1] A. Aghasi and J. Romberg. Convex cardinal shape composition. *SIAM Journal on Imaging Sciences*, 8(4):2887–2950, 2015.

[2] A. Aghasi and J. Romberg. Learning shapes by convex composition. *Under Review, arXiv preprint arXiv:1602.07613*, 2016.

[3] Spencer Aiello, Tom Kraljevic, Petr Maj, and H2O.ai team. *h2o: R Interface for H2O*, 2016. R package version 3.10.0.6.

[4] S. Arora, A. Bhaskara, R. Ge, and T. Ma. Provable bounds for learning some deep representations. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

[5] O. Aslan, X. Zhang, and D. Schuurmans. Convex deep learning via normalized kernels. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, pages 3275–3283, 2014.

[6] F. Bach. Breaking the curse of dimensionality with convex neural networks. *Technical report*, 2014.

[7] Y. Bengio, N. Le Roux, P. Vincent, O. Delalleau, and P. Marcotte. Convex neural networks. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, pages 123–130, 2005.

[8] A. Candel, E. LeDell, V. Parmar, and A. Arora. *Deep Learning with H2O.* H2O.ai Inc., fifth edition, September 2016.

[9] E. Candès and B. Recht. Simple bounds for recovering low-complexity models. *Mathematical Programming*, 141(1-2):577–589, 2013.

[10] A. Choromanska, M. Henaff, M. Mathieu, G.B. Arous, and Y. LeCun. The loss surfaces of multilayer networks. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.

[11] IBM ILOG CPLEX. V12. 1: User's manual for cplex. *International Business Machines Corporation*, 46(53):157, 2009.

[12] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Springer, 2013.

[13] F. Girosi, M. Jones, and T. Poggio. Regularization theory and neural networks architectures. *Neural computation*, 7(2):219–269, 1995.

[14] R. Giryes, G. Sapiro, and A.M. Bronstein. Deep neural networks with random gaussian weights: A universal classification strategy? *IEEE Transactions on Signal Processing*, 64(13):3444–3457, 2016.

[15] Inc. Gurobi Optimization. Gurobi optimizer reference manual, 2015.

[16] K. Hornik, M. Stinchcombe, H. White D. Achlioptas, and F. McSherry. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[17] K. Kawaguchi. Deep learning without poor local minima. In *Preprint*, 2016.

[18] S. Nowlan and G. Hinton. Simplifying neural networks by soft weight-sharing. *Neural computation*, 4(4):473–493, 1992.

[19] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[20] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[21] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[22] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

[23] L. Wan, M. Zeiler, S. Zhang, Y. LeCun, and R. Fergus. Regularization of neural networks using dropconnect. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.